Paul R. MacNeilage, Narayan Ganesan and Dora E. Angelaki

You might find this additional information useful...

This article cites 76 articles, 25 of which you can access free at:
http://jn.physiology.org/cgi/content/full/100/6/2981#BIBL

Updated information and services including high-resolution figures, can be found at:
http://jn.physiology.org/cgi/content/full/100/6/2981

Additional material and information about *Journal of Neurophysiology* can be found at:
http://www.the-aps.org/publications/jn

This information is current as of December 15, 2008 .

Computational Approaches to Spatial Orientation: From Transfer Functions to Dynamic Bayesian Inference

Paul R. MacNeilage, Narayan Ganesan, and Dora E. Angelaki
Dept. of Anatomy and Neurobiology, Washington University School of Medicine, St. Louis, Missouri

MacNeilage PR, Ganesan N, Angelaki DE. Computational approaches to spatial orientation: from transfer functions to dynamic Bayesian inference. *J Neurophysiol* 100: 2981–2996, 2008. First published October 8, 2008; doi:10.1152/jn.90677.2008. Spatial orientation is the sense of body orientation and self-motion relative to the stationary environment, fundamental to normal waking behavior and control of everyday motor actions including eye movements, postural control, and locomotion. The brain achieves spatial orientation by integrating visual, vestibular, and somatosensory signals. Over the past years, considerable progress has been made toward understanding how these signals are processed by the brain using multiple computational approaches that include frequency domain analysis, the concept of internal models, observer theory, Bayesian theory, and Kalman filtering. Here we put these approaches in context by examining the specific questions that can be addressed by each technique and some of the scientific insights that have resulted. We conclude with a recent application of particle filtering, a probabilistic simulation technique that aims to generate the most likely state estimates by incorporating internal models of sensor dynamics and physical laws and noise associated with sensory processing as well as prior knowledge or experience. In this framework, priors for low angular velocity and linear acceleration can explain the phenomena of velocity storage and frequency segregation, both of which have been modeled previously using arbitrary low-pass filtering. How Kalman and particle filters may be implemented by the brain is an emerging field. Unlike past neurophysiological research that has aimed to characterize mean responses of single neurons, investigations of dynamic Bayesian inference should attempt to characterize population activities that constitute probabilistic representations of sensory and prior information.

INTRODUCTION

One of the most challenging and interesting problems faced by the nervous system is reconstructing the state of the world and the state of the organism itself from imperfect sense data. This fundamental problem and its implications have been studied by philosophers and scientists for centuries. In this review, we focus on how the nervous system estimates spatial orientation, which is an organism's sense of body orientation and self-motion relative to the environment. Perceptually, spatial orientation is both immediate and transparent because our sensory, motor, and cognitive systems are highly adapted to this challenge. Behaviorally, spatial orientation is absolutely fundamental to a number of basic motor actions including eye movements, balance, locomotion, and the interaction of any mobile organism with its environment. Computationally, it is a fascinating problem that requires integration of time-varying stimuli from multiple modalities.

Information relevant to spatial orientation comes from diverse sensory systems. Somatosensory and proprioceptive signals from mechanoreceptors in the skin, muscles, and joints can indicate how the weight of the body is supported or moved. However, the modalities that provide the most precise information about spatial orientation are the vestibular and visual senses. Visual information about spatial orientation is extracted from optic flow, the global pattern of visual motion that impinges on the retina during self-motion relative to a stationary environment (Gibson 1950). The resulting visual motion signals are most often proportional to the velocity of the optic flow pattern (Gu et al. 2006). In contrast, the canal and otolith organs of the vestibular system are stimulated by angular and linear accelerations of the head. These dynamic visual and vestibular stimuli are distinct in nature, but the resulting sensory signals are closely related, being approximately integrals or derivatives of each other. Consequently, it makes sense for the nervous system to combine these signals to generate more reliable multi-modal estimates of spatial position, velocity, and acceleration. Exactly how this dynamic multi-modal integration and estimation occurs is the focus of our review.

Neuroscientists have regularly borrowed computational techniques from engineers and mathematicians and applied them to the understanding of biological systems. In the sections that follow, we describe several of these techniques as they have been applied to the problem of estimating spatial orientation: frequency domain analysis, internal models, observer theory, Bayesian estimation, Kalman filtering, and particle filtering. These approaches have been used to investigate a wide variety of spatial orientation behaviors. Some have been used to characterize reflexive eye movements (e.g., Angelaki et al. 1999; Green and Angelaki 2004; Green et al. 2007; Raphan et al. 1977; Robinson 1981; Merfeld et al. 1993), others perception of body orientation and self-motion (e.g., Borah et al. 1988; Bortolami et al. 2006; Laurens and Droulez 2007; MacNeilage et al. 2007; Zupan et al. 2002), others the onset of motion sickness (e.g., Bos and Bles 2002; Oman 1982), and still others postural responses (e.g., Kuo 2005; van der Kooij et al. 2001). We refer to several of these models here because they address the same problem; some form of spatial orientation estimate is necessary for perception as well as control of reflexive eye movements and postural responses. Note that different spatial orientation estimates may underlie perception and motor action (Merfeld et al. 2005; Zupan and Merfeld 2005).

In addition to examining the advantages and drawbacks of each technique, we also discuss the levels of analysis at which each technique operates. Marr (1982) proposed three levels of analysis for the investigation and understanding of computational processes. The level of *computational theory* describes the overarching goal and logic of the computation. For example, a project that analyzes spatial orientation at the level of computational theory might seek to determine if the nervous system combines information probabilistically in a way that maximizes the chances of estimating the true body orientation and self-motion of the observer. The level of *representation and algorithm* describes more specifically how the computation may proceed, and how information may be represented and processed from input to output. An example of a project operating at this level would be one that aims to determine what coordinate system sensory signals may be coded in or one that investigates the algorithm the nervous system may use to disambiguate gravitoinertial force. The level of *hardware implementation* describes how the system could be realized physically. Investigations at this level seek to determine how populations of neurons in different areas of the brain could actually execute the algorithms necessary to achieve spatial orientation. This hierarchical conceptualization of computational processes emphasizes the independence of each level from the one above it; a given computational goal can be achieved by a number of alternative algorithms, and a given algorithm may be neurally implemented a number of ways.

Many of the techniques we will describe operate at one or two of these levels; this goes some way toward defining their strengths and weaknesses. However, we emphasize, as Marr (1982) did, that full comprehension is not possible without some amount of understanding at all three levels. The aim of this review is to provide an overview of current approaches to estimation of spatial orientation that might contribute to a more comprehensive and integrated understanding for researchers working at all three levels of analysis.

Velocity storage: an example for comparison

We have defined spatial orientation in the broadest terms to encompass a wide range of phenomena. However, because our goal is to compare different computational approaches, we have chosen to focus on velocity storage, an aspect of spatial orientation that has been modeled using virtually every approach. Velocity storage describes the observation that the angular velocity estimate of the CNS (as expressed by reflexive eye movements and perceptual measures) outlasts the angular velocity signaled by the afferent signals from the semicircular canals (see Fig. 1).

Historically, this phenomenon has played a central role in investigations of spatial orientation processing because it is representative of the challenges inherent in dynamic state estimation problems. The physical stimulus to the semicircular canals is angular acceleration of the head, but the quantities of interest are angular velocity and position (i.e., head tilt). The velocity estimate may be used to generate motor commands that drive reflexive eye movements for gaze stabilization (see Raphan and Cohen 2002 ; Robinson 1981 for reviews). The position estimate provides information about head tilt that is vital for interpreting ambiguous otolith signals (Green and Angelaki 2003, 2004; Green et al. 2005; Merfeld 1995; Mer-
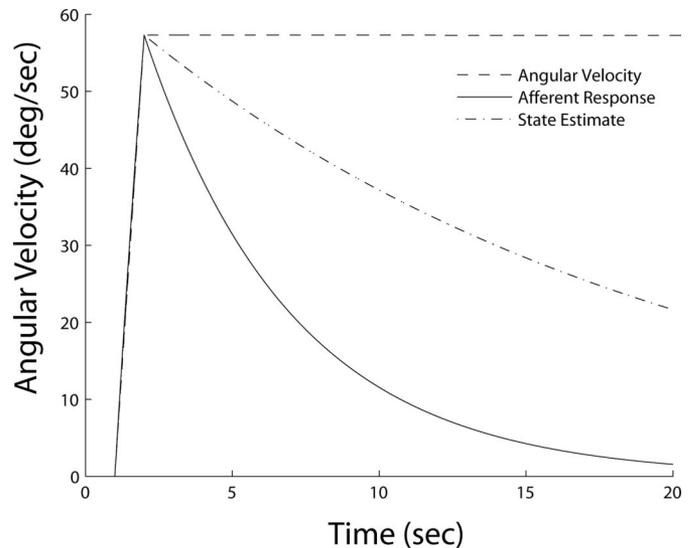

FIG. 1. Schematic illustration of velocity storage. For a constant velocity stimulus (- - -), the angular velocity estimate of the CNS (-·····) outlasts the angular velocity signaled by the afferent signals from the semicircular canals (—).

feld et al. 1993; Yakusheva et al. 2007; Zupan et al. 2002). Thus investigations of velocity storage are primarily concerned with the mechanics and functional significance of these integration processes.

As discussed below, the functional significance assigned to the lengthening of the canal time constant is tightly coupled to the particular computational approach. For example, velocity storage modeling started with frequency domain analysis (see next section). These models assume that a central low-pass filter *stores* angular velocity signals from the semicircular canals to lengthen the time constant of the vestibuloocular reflex (VOR) and maintain a more accurate estimate of angular velocity during constant velocity rotations (Cohen et al. 1977; Raphan et al. 1979). Alternatively, internal model and observer theory approaches assume this lengthening is due to a low-pass filter within the internal model of sensor dynamics (Merfeld 1995; Merfeld et al. 1993). Finally, probabilistic methods, like Kalman and particle filtering, suggest that the velocity storage time constant depends on noise that accumulates during the integration process. We review each of these techniques in turn starting with classical frequency domain analysis.

Transfer function approach: classical frequency domain analysis for linear dynamic systems

The transfer function approach provides a way to represent and analyze the input-output relationships of linear dynamic systems in the frequency domain. It has been used for some time and has been extensively described and evaluated elsewhere (Gaskill 1978; Robinson 1981). Here we briefly review how this approach has been used to investigate velocity storage. As described above, investigations of velocity storage are typically concerned with how the vestibular system integrates the angular acceleration stimulus to obtain estimates of velocity and position. The first integration from acceleration to velocity is achieved primarily by the physical dynamics of the semicircular canals. In the frequency domain, the dynamics of the canals for a certain range of head velocities (H) may be

described by a first-order dynamic system and the following transfer function

$$\frac{\Delta R_c(s)}{\dot{H}(s)} = \frac{s}{s + T_c} \qquad (1)$$

where R_c is the afferent response of the canals, T_c is the time constant, and s is the complex frequency (Goldberg and Fernandez 1971). Notice that the angular velocity signaled by the canals is initially accurate. However, for a constant velocity stimulus, the response decays toward zero, becoming less and less accurate over time (Fig. 1, —).

The dynamics of the rotational vestibuloocular reflex may be described by a similar transfer function

$$\frac{\dot{E}(s)}{\dot{H}(s)} = \frac{s}{s + T_{vor}} \qquad (2)$$

where $\dot{E}$ is eye velocity and T_{vor} is the time constant of the VOR. To a first approximation, the primary difference between the dynamics of the canal afferent signal (*Eq. 1*) and the reflexive eye movements (*Eq. 2*) is the lengthening of the time constant ($T_c < T_{vor}$). This lengthening means that the signal decays more slowly and thus remains more accurate for a longer period of time during low-frequency stimulation (Fig. 1, -·-). Two model types were initially proposed to account for the lengthening of the VOR time constant (velocity storage), one using a low-pass filter and positive feedback (Robinson 1977) and the other a combination of direct and low-pass filter pathways (Raphan et al. 1977, 1979).

The transfer function approach has allowed researchers to precisely describe how sensory input and motor output may be represented in terms of neural activity (e.g., Galiana and Outerbridge 1984; Green and Galiana 1998; Green et al. 2007). These representations, in turn, have placed constraints on the processing that must occur in intermediate stages, which has served to guide further investigations. This approach has led to great advances in our understanding of the neural processing underlying velocity storage (Waespe and Henn 1977a,b, 1979, 1981; Waespe et al. 1981) and oculomotor function more generally. For example, Skavenski and Robinson (1973) used this technique to analyze the VOR dynamics over a wide range of frequencies. They discovered a phase lag between head position and eye position that could not be accounted for by the dynamics of the canals or oculomotor plant. Based on this observation, they concluded that there must be a neural network in the VOR pathway that integrates velocity into an eye position signal. Subsequent research has probed the responses of neurons in the brain stem and cerebellum in an attempt to isolate this neural integrator (Fukushima 1991; Goldman et al. 2002; McFarland and Fuchs 1992; Robinson and Fuchs 2001; Scudder and Fuchs 1992). As a consequence, we now have a much better understanding of the neural processes underlying gaze stabilization, and how the processing differs depending on the whether the stimulus is visual or vestibular, rotation or translation, and actively or passively generated (reviewed in: Angelaki 2004; Angelaki and Cullen 2008; Cullen and Roy 2004).

However, in isolation, the transfer function approach has several shortcomings. In terms of level of analysis, it operates primarily at the levels of hardware implementation and representation. The transfer function simply connects the input to the output. It does not address the question of *why* the time constant of velocity storage is lengthened or the computational goals that are served by this low-pass filtering. Furthermore, to characterize neural processing with a single transfer function, one must assume that the activity of a pool of neurons can be approximated by the typical neuron (Robinson 1981). Variability in neural responses, also known as noise, is thus ignored. Next we describe how the transfer function approach has been expanded to more general concepts of linear systems analysis, in particular those related to internal models and observer theory. Notice that although the concepts of internal models and observer theory still assume deterministic computations and largely ignore noise and probabilistic representations, they nevertheless constitute a valuable advance in applying linear systems analysis tools to biological processing.

Internal models: neural representations of physical principles

Internal models suggest that the nervous system maintains internal representations of certain physical variables that are essential to sensory and motor processing. They go beyond simple input-output representations to generate hypotheses about the internal processing and algorithms that underlie the behavior of the system as a whole (for reviews, see Tin and Poon 2005; Wolpert and Kawato 1998). In essence, internal models form the conceptual link between input/output transfer function representations (previous section) and modeling complex simultaneous interactions of multiple dynamic variables (observer theory, next section).

The concept of internal models is quite general, but specific classes of internal models have been proposed in particular contexts (Fig. 2). In motor control, *inverse motor models* are required to map desired movement into motor commands to control muscle movement, and this mapping depends on the dynamics of the motor plant. For example, the neural integrator example described earlier (Skavenski and Robinson 1973) fits conceptually with the notion of an inverse model of the oculomotor plant. It is responsible for converting desired movement into motor commands in response to various kinds of visual and vestibular stimuli (cf. Ghasia et al. 2008; Glasauer 2007; Green et al. 2007; Robinson 1981). *Forward motor models* perform the same transformation in reverse; they convert efference copies of motor commands into internal estimates of executed movement. Such estimates of executed movement may be compared with sensed movement to verify that the plant is functioning properly. This is an important mechanism for maintaining calibration and robust motor control even when the dynamics of the plant change, as during development, injury, or aging (cf. Wolpert and Kawato 1998).

For sensory processing, internal models of *inverse sensor dynamics* allow the nervous system to convert afferent sensory signals into estimates of the sensory stimuli via a mathematical representation of the sensor. For example, the nervous system could estimate angular velocity from the canal afferent signal using a sensor model, i.e., something like the inverse of the canal transfer function expressed in *Eq. 1*. An inverse model of sensor dynamics attempts to reconstruct the physical stimulus from dynamically deficient afferent activity. Such an inverse model for sensor dynamics could provide a conceptual moti-
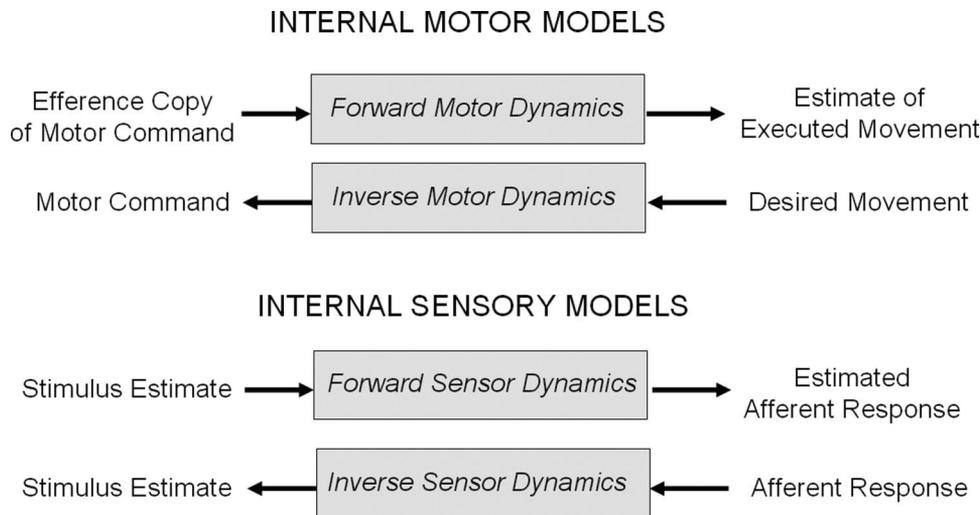
INTERNAL MOTOR MODELS

Efference Copy of Motor Command → [Forward Motor Dynamics] → Estimate of Executed Movement

Motor Command ← [Inverse Motor Dynamics] ← Desired Movement

FIG. 2. Examples of internal models for sensory and motor processing.

INTERNAL SENSORY MODELS

Stimulus Estimate → [Forward Sensor Dynamics] → Estimated Afferent Response

Stimulus Estimate ← [Inverse Sensor Dynamics] ← Afferent Response

vation for the combination of direct and low-pass filter pathways proposed as the mechanism for lengthening the VOR time constant by Raphan and colleagues (1979). A model of *forward sensor dynamics* calculates the expected afferent signal for a given sensory stimulus. The expected afferent signal may then be compared with the actual afferent signal to verify that the sensory system is operating normally. The functional significance of the forward sensor model parallels that of the forward motor model, to maintain calibration and performance even in the face of changing system dynamics (cf. Wolpert and Kawato 1998).

Notice that within the framework of internal models, the two examples of low-pass filtering and direct/indirect pathways proposed by Robinson and Raphan for the oculomotor neural integrator and the velocity storage integrator, respectively, now have a deeper conceptual meaning; each can be thought as the inverse model of the motor plant and sensor dynamics, respectively. However, the forward model concept is more recent and, in many respects, less intuitive than the inverse dynamic model. As will be developed in the next section, forward models represent critical components of any observer-like control system. But before we move on to describing the observer theory framework, there is one additional internal model concept that is of relevance for spatial orientation.

In addition to sensory and motor dynamics, internal models of *physical laws* have been proposed. For example, body orientation is estimated by sensing gravitational force (cf. Mayne 1974). However, gravitational force and inertial force due to linear acceleration are indistinguishable and cannot be encoded separately. Only their sum can be sensed (Angelaki et al. 2004; Fernandez and Goldberg 1976; Si et al. 1997). This can be expressed as

$$f = g - a \tag{3}$$

where g is gravitational force, a is inertial force and f is the summed gravitoinertial force (GIF; all are vectors). To estimate gravitational force in the presence of inertial force, the nervous system must have an internal model of this physical relationship. In addition, the nervous system can use knowledge about how the direction of gravity in the head changes during head rotations, which can be expressed as follows

$$\dot{g} = -\omega \times g \tag{4}$$

where ω is the angular velocity of the head and $\dot{g}$ is the time derivative of gravitational force. This constitutes another internal model of physical quantities. There are several aspects of *Eq. 4* that are worth emphasizing. First, it represents a vector differential equation that is highly nonlinear (see Green et al. 2005 for a further explanation of this computation). Thus, unlike the simple example of lengthening of the canal time constant (velocity storage), the problem of spatial orientation is in general nonlinear. Second, the concept of internal models of physical laws can be thought of as *an expanded form of sensor internal models*. Remember that sensor internal models relate the sensory response to the physical quantities of the outside world. In the case of the semicircular canal system, at least to a first approximation, this relationship is simple and described by a low-pass filter. In the case of the otolith system, however, this relationship is not only temporal. As a result, a simple temporal filter is not sufficient to reconstruct inertial acceleration from otolith afferent activity. Instead, there is a complex, highly nonlinear, spatiotemporal version of the forward and inverse models associated with the otolith system, as described by *Eqs. 3* and *4*.

Several studies (Angelaki et al. 1999, 2004; Merfeld et al. 1999, 2005) have measured eye movements, perception, and neural responses and obtained results consistent with the notion that the nervous system uses such internal models to estimate spatial orientation from highly deficient and ambiguous otolith afferent activity. Perhaps more important is the demonstration that although central canal responses are well-behaved and only temporally transformed relative to the semicircular canal properties, brain stem, thalamic, and cerebellar neurons exhibit translation responses with strong spatiotemporal dependencies and highly nonlinear, gravity-dependent response properties (Angelaki and Dickman 2000; Dickman and Angelaki 2002; Meng et al. 2007; Shaikh et al. 2005; Yakusheva et al. 2007). While keeping in mind that spatial orientation processing is in general nonlinear, we return to the simple example of velocity storage in the following sections, because our primary aim is to illustrate the various computational approaches, not to explore the complexities of the spatial orientation problem itself.

Observer theory: state estimation via error monitoring

As described in the preceding text, internal models propose the existence of internal representations that may not be directly observable as sensory input, or motor output. These internal quantities can be thought of as *states* of the system, and observer theory proposes a technique for estimating the states of complex, dynamic systems. A single observer model can have multiple states and can therefore incorporate many of the internal models described in the preceding text. Thus observer theory provides a way to model complex simultaneous interactions of multiple dynamic variables in a unified framework.

A generic sensorimotor observer is depicted in Fig. 3; these models are known as "observers" because they monitor or observe only the input, $u(t)$, and output, $y(t)$, but they infer internal states of the system. Everything within the upper dashed line is part of the *system*, and everything within the lower dashed line is part of the *observer model* of the system. The observer functions by comparing the output and the estimated output to generate a feedback signal that is passed through an internal model of motor dynamics to generate the state estimate. The state estimate, in turn, is passed through the model of sensor dynamics to generate the estimated output, which is used to generate a feedback signal, and so on.

Observer models originated in engineering as a method for controlling system dynamics; by monitoring the input and output an observer can estimate the states of the system and use those estimates to generate a control signal to drive the system to the desired state. Merfeld et al. (1993) proposed a simple passive observer model for estimating yaw angular velocity of the head in response to canal afferent signals. The model is passive because it does not compare estimated and desired states to generate a control signal (omit gray lines on the left of Fig. 3); the input, $u(t)$, is set to zero and motor dynamics are not relevant (i.e., set to unity transfer). Passive observers are suited to modeling passive perception and estimation because they only function to estimate the state of the system that is only affected by the passive stimulus. The stimulus (angular velocity) is passed through the canal sensor dynamics to generate an afferent signal, $y(t)$. The observer calculates a feedback signal. This feedback signal is multiplied by a gain to generate an estimate of the state of the system, i.e., an estimate of the state of the semicircular canals and the angular velocity stimulus. The estimated state is then passed through the forward model of sensor dynamics to generate the estimated sensory afferent signal, $\hat{y}(t)$.

This simple arrangement ensures that the estimated angular velocity (the state) is equal to the change in firing rate caused by the passive stimulus (the feedback signal) times the gain; the time constant for the angular velocity estimate therefore depends only on what value is chosen for the gain (see *Eq. 5* in Merfeld et al. 1993). This is an unconventional observer in the engineering sense because the state estimate does not converge to the true state. Instead for a constant velocity input, the estimated state converges to zero angular velocity. Their observer model achieves velocity storage via negative feedback and a forward internal model of canal dynamics that differs from the velocity storage models initially proposed by Robinson (1977) (low-pass filter, positive feedback) and Raphan et al. (1977) (parallel direct/indirect paths).

An important advantage of the observer approach is that multiple internal models may be incorporated into a single observer such that any number of states can be monitored and different state estimates can interact. Oman (1982), Merfeld (1995), Bos and Bles (1998, 2002), and Zupan et al. (2002) propose more comprehensive observer-like models for spatial orientation that incorporate internal models of both semicircular canal and otolith dynamics and also internal models of physical laws equivalent to those shown in *Eqs. 3* and *4* (see Bos and Bles 2002 for a review).

Several of these comprehensive observers have come to be known as *sensory conflict models for GIF resolution* because they aim to resolve sensed GIF into separate estimates of gravitational and inertial force; the models generate internal state estimates of gravity, linear acceleration and angular velocity. As with the simple observer model for velocity storage described in the preceding text, the mechanism that drives this resolution process is detection of conflicts, or
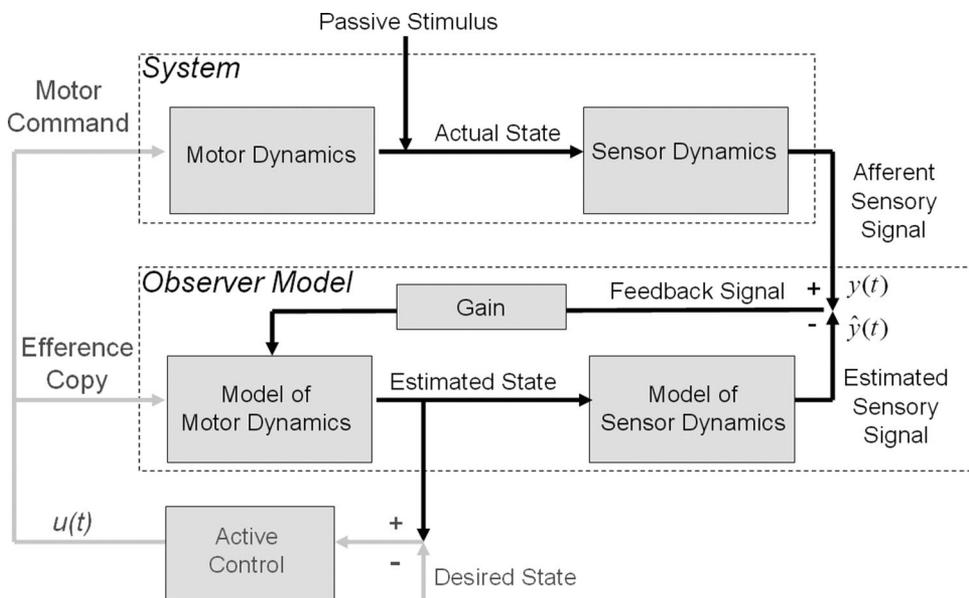

FIG. 3. Traditional observer model architecture (after Merfeld et al. 1993). The observer functions by comparing the output and the estimated output to generate a feedback signal that is passed through an internal model of motor dynamics to generate the state estimate. The state estimate, in turn, is passed through the model of sensor dynamics to generate the estimated output, which is used to generate a feedback signal, and so on. For passive observers, the input, $u(t)$, is set to 0, motor dynamics are not relevant (i.e., set to unity transfer), and there is no comparison between estimated and desired states to generate a control signal (omit gray lines). Passive observers are suited to modeling passive perception.

mismatches, between expected and observed sensory signals. However, in these more comprehensive models, several sensory conflict signals are detected and these are multiplied by several different gains before they are fed back into the system (see original texts for more details).

The GIF resolution observer succeeds in characterizing both perception and reflexive eye movements in response to complex spatial orientation stimuli under a variety of circumstances (Merfeld 1995; Merfeld et al. 1993). However, the explanatory power of such models is limited because certain features of the model architecture and parameters are somewhat arbitrary. For example, there is no a priori rationale for how the feedback gains should be set. The gains are free parameters and thus the dynamics of the system are simply fine tuned by the modeler to agree with empirical observations. In addition, the state estimates are performed in the absence of noise. Yet most if not all physical systems and real world signals are noisy, so it is imperative to consider noise or uncertainty during state estimation. Kalman filtering provides a probabilistic alternative to observer modeling. However, before describing the Kalman modeling approach, we will first briefly review Bayesian theory.

Bayesian theory: statistically optimal estimation in the presence of uncertainty

The observer models described above are said to be deterministic because uncertainty or noise associated with the input and output signals is ignored, just as it was in the transfer function approach. In reality, sensory input and motor output signals are stochastic or noisy. Noise affects performance of the system and should therefore be included in the model.

Bayesian theory describes how statistical inferences should be made in the presence of noise or uncertainty. It is governed by Bayes' rule

$$p(x|y) = \frac{p(y|x)p(x)}{p(y)} \qquad (5)$$

The term on the left is the posterior probability distribution and it expresses the probability associated with each possible state (x) given the information (y) that is available. The first term in the numerator on the right, $p(y|x)$, is the likelihood distribution which expresses the probability of the information given the state; for sensorimotor processing, this is functionally equivalent to a forward model of sensor dynamics. The second term in the numerator on the right, $p(x)$, is the prior probability distribution which expresses the probability of the state, even in the absence of any information. The term in the denominator, $p(y)$, is the probability of the information, but its importance is minimal as it simply serves to scale the probability distribution expressed in the numerator.

Bayesian theory has been used to model perception in a variety of circumstances. In these perceptual models, the likelihood distribution represents the sensory information, the prior probability distribution represents preexisting knowledge or experience, and the posterior probability distribution represents the information that determines the percept. The peak of the posterior distribution is the most probable state of the world given the sensory and prior information available to the organism.

Bayesian models have successfully characterized perceptual responses under various conditions. Illusory visual motion percepts were found to be well predicted by a Bayesian model that includes a prior for slow motions (Weiss et al. 2002). Bayesian models are also well-suited to modeling/explaining the integration or combination of redundant sensory estimates whether they arise from a single modality (e.g., Hillis et al. 2004) or different modalities (e.g., Ernst and Banks 2002; Ghahramani et al. 1997); given the assumption of independent (uncorrelated) noise, each sensory signal is weighted based on its reliability (or inverse variability).

This approach has been applied recently to perception of spatial orientation. MacNeilage et al. (2007) proposed a Bayesian model for disambiguation of gravitoinertial force by visual cues to body orientation (e.g., horizon) and self-motion (e.g., accelerating optic flow). They found that visual cues were used to interpret ambiguous vestibular cues and that the weight given to these cues was less when the cues were less reliable, consistent with predictions of the Bayesian model. They proposed two priors for perception of spatial orientation, one for zero inertial acceleration, and another for upright body orientation, also suggested by Eggert (1998); this prior is functionally similar to the idiotropic vector proposed by Mittelstaedt (1983). de Vrijer, Medendorp, and van Gisbergen (2008) subsequently showed that a Bayesian prior for upright orientation can account for consistent underestimation of the visual vertical at large body tilt angles (Aubert effect); they showed that the underestimation is proportional to the variability of the body tilt estimate, as predicted by a Bayesian model.

Despite success characterizing perception, Bayesian models have several limitations. They operate only at the level of computational theory; they do not explicitly consider how the sensors convert the sensory stimulus into a neural representation and how these neural representations may be processed prior to integration. In addition, such Bayesian models are usually limited to discrete or static situations. In actuality, real-world stimuli and the associated percepts and motor actions are usually dynamic and continuous. Under these circumstances, observer theory and the Bayesian framework can be combined in the form of Kalman filtering.

Kalman filtering: statistically optimal state estimation for linear dynamic systems with noise

Kalman filters combine the strengths of observer theory and Bayesian approaches. They are observer models, and therefore capable of modeling complex, dynamic, continuous systems. Kalman filters have essentially the same structure as that illustrated in Fig. 3. Actual and expected signals are compared and fed back into the system with some gain. However, unlike observer models, the gain applied to the feedback signal is not arbitrary; it is determined in a statistically optimal (Bayesian) fashion that depends on the noise or probability distributions associated with the sensory and motor signals. Thus a Kalman filter can be thought of as a Bayes-optimal observer model for noisy dynamic systems.

The key to implementing a Kalman filter lies in computing the statistically optimal Kalman gain. To compute this gain analytically, we must convert the expression of system dynamics from the transfer function representation into the state space representation (see APPENDIX A for details). In the state space

representation, the states are represented by the state vector, x. Note that the state variables making up this vector need not correspond directly to quantities of interest, but importantly, any dynamic quantity of the system can be expressed as a weighted linear combination of the state variables. The size of the state vector depends on the order of the differential equation that describes the dynamics of the system.

The state space representation requires rewriting the dynamics of the system by a first order vector differential equation that quantifies $\dot{x}$, the time rate of change of the state of the system at every instant, as follows

$$\dot{x} = Ax + Bu(t) \qquad (6)$$

$$y(t) = Cx \qquad (7)$$

where $y(t)$ and $u(t)$ are the values of the output and input at time t. A, B, and C are matrices the values of which depend on the coefficients of the numerator and denominator polynomials of the transfer function (see APPENDIX A). A is the system matrix that describes the main dynamics of the system, B describes the input dependence of the rate of change of state, and C describes the state dependence of the output. APPENDIX A shows how to calculate these matrices for a single-input, single-output system.

To understand how the optimal Kalman gain depends linearly on the noise in the system, consider the following state space equations

$$X_{k+1} = A_k X_k + b_k u_k + \xi_k \qquad (8)$$

$$y_k = C_k X_k + \theta_k \qquad (9)$$

These equations function to predict the state at time $k + 1$ given the information at time k. Notice that they are essentially discrete expressions of *Eqs. 6* and *7* with noise added. Parameters ξ_k and θ_k are zero mean normally distributed (Gaussian) random variables representing input (i.e., process) and output (i.e., measurement) noise, respectively. These noises are assumed to be uncorrelated.

As can be seen, the states (X_k) and the output (y_k) are also normally distributed random variables (vectors) that are causally and linearly correlated to the noises ξ_k and θ_k. For exam-

ple, the random variable y_k is correlated to the state X_k via the output equation (*9*) and hence to the noise $\{\xi_1, \ldots, \xi_k\}$ and θ_k. In this case, state estimation can be recast as a problem of finding the posterior probability distribution of the state at time instant k in terms of the estimate of state at the previous time instant ($k - 1$) and the likelihood of the current output y_k given the state (Berger 1985). Bayes' rule implies that

$$p(X_k|y_k) \propto p(y_k|X_k) \int_{X_{k-1}} p(X_k|X_{k-1})p(X_{k-1})] \qquad (10)$$

The most likely state at time k is the peak, or expected value, of this posterior distribution. The posterior distribution can be thought of as being statistically optimal because the above relation utilizes all the available information to come up with an estimate according to the Bayesian principle. APPENDIX B shows the implementation of Kalman filters for discrete time systems, and the generalization to continuous time systems.

It is interesting to note that static Bayesian and dynamic Kalman filter models (as described so far) achieve optimality in different ways: static models combine estimates from different sources, whereas Kalman filters combine estimates over time. However, Kalman filters are also well-suited to modeling optimal combination across modalities because they operate on vector ordinary differential equations.

Kalman filters have often been used to model state estimation problems faced by the nervous system (Todorov 2004; Wolpert and Gharamani 2000). Statistically optimal multimodal estimation of spatial orientation by the human nervous system was proposed early on by Young (1970). Later, Borah et al. (1988) implemented a Kalman filter version of the passive observer model illustrated in Fig. 4, which compares expected to observed canal afferent signals to estimate angular velocity. While the overall structure of their Kalman filter is quite similar to the observer model of Merfeld et al. (1993), there are a couple of significant differences in the implementation. These differences were probably introduced to make the state space computations more convenient.

First, they consider the input to the canals to be angular acceleration, not velocity, which makes the canal transfer
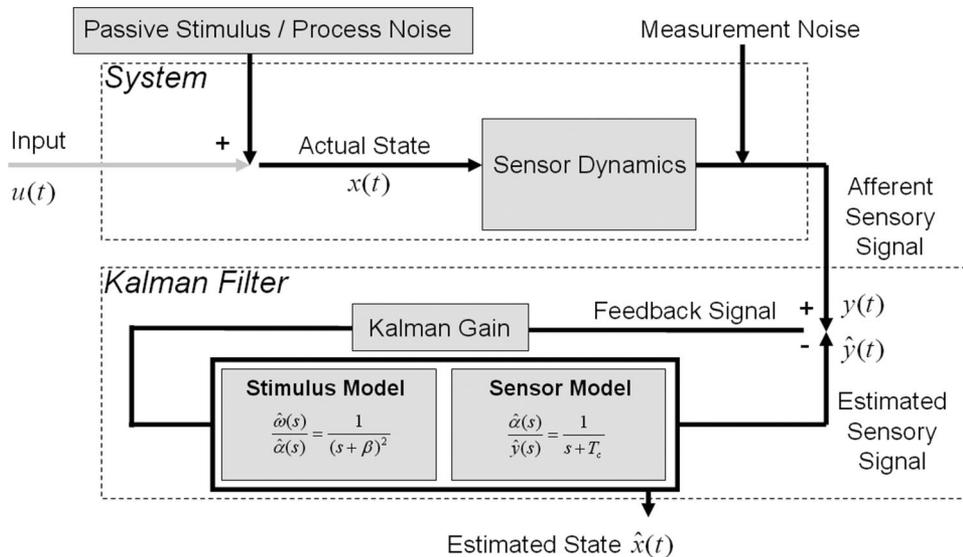

FIG. 4. Kalman filter model (after Borah et al. 1988). A simplified version is depicted to illustrate how the Kalman filter approach parallels the Observer model in Fig. 3. The model resembles a passive observer, so input is set to 0, and motor dynamics (not shown) are set to unity transfer. The feedback signal is multiplied by the statistically optimal Kalman gain. The result is passed through the internal model of the system dynamics, which includes the stimulus internal model and the internal model of sensor dynamics. Note that the Kalman filter uses the state space representation of these internal models.

function a proper rational function (one in which the order of the numerator is less than that of the denominator), thereby guaranteeing well-behaved state space equations (see APPENDIX A). Second, to specify the relationship between the state (x) and the rate of change of state ($\dot{x}$) in the state space equations, they include a *stimulus internal model* that relates estimated angular velocity and acceleration to the rate of change of acceleration

$$\frac{\hat{\omega}(s)}{\hat{\alpha}(s)} = \frac{1}{(s + \beta)^2} \qquad (11)$$

thereby closing the dynamics of the system. The observer model implemented by Merfeld et al. (1993) avoids these complications because the state space representation is not required for observer models as it is for Kalman filters.

Figure 5 shows the response of a Kalman filter model like that depicted in Fig. 4 to a step in yaw angular velocity (red dashed line). The model makes a very interesting prediction, namely that the time constant for velocity storage depends on the noise in the system. However, the time constant also depends on an artificial parameter of the model, β, the bandwidth cutoff of the stimulus internal model.

As with the observer model described in the preceding text and because of the identical negative feedback structures of the two models, the Kalman gain determines the time constant and gain of the state estimate of angular velocity, $\hat{x}(t)$. However, the difference is that in this case the gain is not arbitrary. The statistically optimal Kalman gain is determined by the Ricatti Equation (APPENDIX B), which uses the state space representation of the stimulus internal model and the sensor model. Therefore the gain is influenced by both parameters, namely the noise on the canal signal and the bandwidth cutoff of the stimulus internal model. On these grounds, it could be argued that the Kalman filter is a step backwards from the observer model because it replaces one free parameter (gain) with two

(noise and bandwidth cutoff). Indeed, the Borah et al. model is a chronological precursor to the Merfeld model. However, the advantage of the Kalman filter approach, and our motivation for presenting it after the observer model, is that the free parameters are conceptually motivated; they represent noise in the system and some form of frequency dependent prior expectations about vestibular stimuli. The fact that the parameters correspond to real-world quantities like noise, that may be empirically measured, means that the Kalman filter model can be tested and has the potential to predict and possibly explain characteristics of behavior that cannot be explained by post hoc parameter fitting.

Other examples of Kalman filtering applied to the problem of spatial orientation come from the field of postural control, which requires muscle contractions and thus motor commands in response to vestibular and other sensory information. These postural control models are more similar to the active observer model depicted in Fig. 3, in which estimated and desired states are compared to generate a control signal. For example, Kuo (1995, 2005) proposed a model that takes motor command signals to leg muscles as input, treats visual and vestibular estimates of body motion as output, and generates state estimates of body movement and leg joint configuration. The covariance of the ankle versus hip sway was the only observable data that could be experimentally verified against that predicted by the model. The data predicted by the model were corroborated against different physical/experimental scenarios that involved independent and conflicting visual and vestibular inputs, corresponding to different conditions governing inaccurate body sensors via a view screen and a motion platform. The data were also tested against and verified for different sensory conditions and different age groups and were found to corroborate well with experimental data. However, because Kalman filters are only valid for linear systems, Kuo (2005)
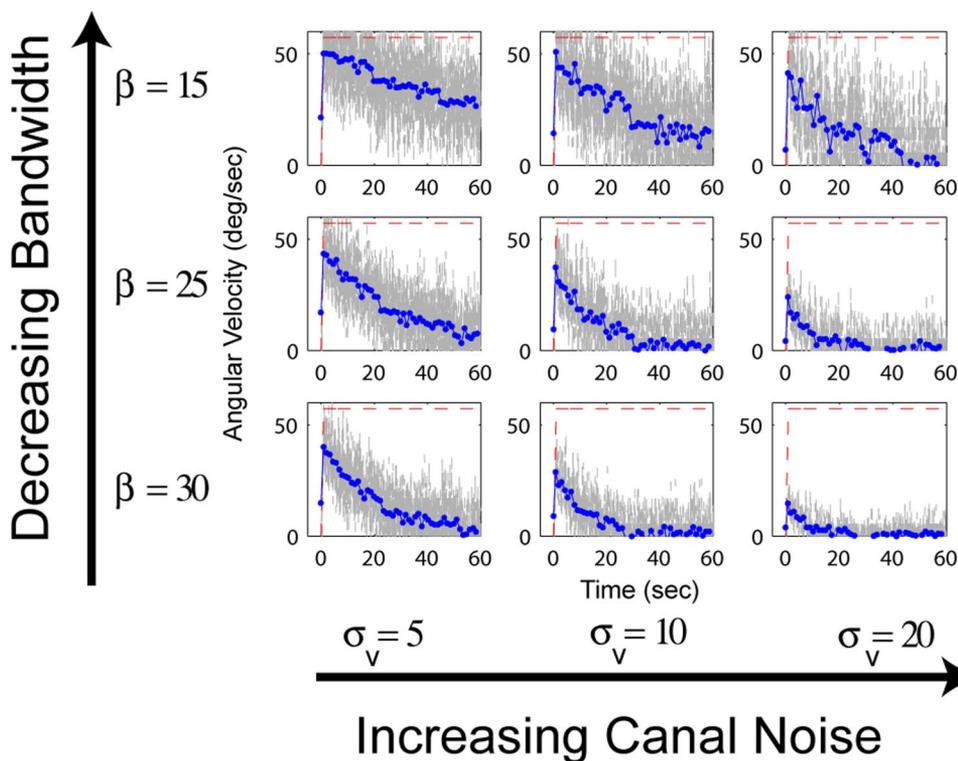

FIG. 5. In a Kalman filter model, the time constant of the angular velocity estimate (velocity storage) depends on 2 parameters: the SD of the Gaussian-distributed noise on the canal signal (σ_v) and the bandwidth cutoff of stimulus internal model (β). Values of canal noise increase from left to right from 5 to 20°/s. Values for the bandwidth cutoff increase from top to bottom. The red dashed lines indicate the actual angular velocity, the gray lines show the estimated angular velocity, and the blue lines show the mean of that estimate during each second of the simulation. A 1st-order approximation of the transfer function proposed by Borah et al. (1988) was used; this does not affect behavior of the model for the frequencies of stimulation examined here.

restricted his analysis to small deviations from upright. In addition, noise values for input and output signals were treated as free parameter and used to fit model behavior to empirical observations.

Another Kalman filter model of postural control was proposed by van der Kooij and colleagues (1999, 2001). The implementation of this model is similar to that proposed by Kuo except this Kalman filter is modified to deal with system nonlinearities. The mechanism that allows this extension to nonlinear systems is a cascade combination of an extended Kalman filter and a nonlinear predictor. This scheme is inherently suboptimal because it uses linear equations for prediction and estimation even though the motor dynamics being modeled are nonlinear. As a consequence, any inaccuracies in the internal model of the dynamics or the knowledge of noise parameters will cause the filter to quickly diverge or perform poorly. Despite these limitations, extended Kalman filters are easily implemented and therefore widely used in discrete time inertial navigation systems and mobile robots where position sensing in the environment is one of the major problems (Awalt et al. 2003; Spingarn 1987). Indeed, exchange of ideas between the designers of such systems and researchers investigating human cognition could be beneficial to both communities.

Noise in sensory and motor signals is the fundamental constraint on sensorimotor performance, so sensorimotor processing cannot be understood without considering uncertainty (Faisal et al. 2008; Wolpert 2007). The Kalman filter approach is promising because it provides a way to incorporate the uncertainty or noise inherent in all physical systems. However, Kalman filter models are rigid and inflexible in some respects and therefore not optimally suited to all estimation and control problems. For example, they are best suited to linear systems; modifications of the Kalman filter to deal with nonlinear systems (e.g., adaptive Kalman filter of van der Kooij et al. 2001) are limited and suboptimal. This is problematic because nonlinear computations (e.g., resolving gravitoinertial force) are fundamental to maintaining spatial orientation. In addition, Kalman filters are restricted to assume Gaussian-distributed noise. While this assumption may be valid in many cases, it is not always justified (cf. Faisal et al. 2008). Finally, the feedback structure of Kalman filters is quite rigid. Inherent in any diagram of Kalman filter architecture are explicit assumptions or hypotheses about which signals are being compared to generate the feedback signals and how those feedback signals are passed back to the system. In the next section, we introduce a more distributed and flexible probabilistic modeling technique that does not assume specific feedback circuits, Gaussian noise, or linear system dynamics.

Particle filtering: a simulation technique for dynamic Bayesian state estimation

Particle filtering is a method for modeling dynamic Bayesian inference that is purely probabilistic in that all information is represented by full probability distributions. It is a simulation technique in which these probability distributions are sampled at each time step. The samples are known as "particles"; they are run through the system dynamics, subject to stochastic processes (noise), and the results are used to recreate the probability distribution at the next time step. Conceptually it is convenient to think of the particle filter as a recursive version of the static Bayesian models described previously, such that the process of Bayesian inference is extended over time. The method relies heavily on Bayes' rule (*Eq. 5*), which expresses the posterior distribution, $p(x|y)$, as the product of the likelihood distribution, $p(y|x)$ and the prior, $p(x)$, where x is the state and y is the sensory information (see Bayesian section for more detail). In the particle filter, the posterior is expressed as $p(x_t|y_{1:t})$, such that current state, x_t, depends on all previous sensory information, $y_{1:t}$. The instantaneous likelihood is expressed $p(y_t|x_t)$, and the prior is expressed as $p(x_t|x_{t-1})$ and thus describes the dependence of the current state on the previous state. APPENDIX C shows how the particle filter posterior can be calculated given the particle filter likelihood and prior distributions.

Laurens and Droulez (2007) have used the particle filtering method to develop a dynamic Bayesian inference model for the perception of spatial orientation from vestibular sensory information. The problem was cast as that of estimating the three-dimensional (3D) state of the head given vestibular sensory and prior information. The state of the head, x, is a vector that describes the 3 df for angular position, Θ, angular velocity, Ω, deflection of the cupula, C, and linear acceleration, A, for a total of 12 state variables. The sensory information, y, describes the 3 df for afferent signals from the semicircular canals, V, and otoliths, F, for a total of six output variables. So, $x_t = (\Theta_t, \Omega_t, A_t, C_t)$ and $y_t = (V_t, F_t)$.

The authors calculate the particle filter likelihood distribution as follows

$$p(y_t|x_t) = p(V_t|C_t)p(F_t|\Theta_t, A_t) \qquad (12)$$

The first term on the right says that the canal afferent response depends on cupular deflection. The authors assume that canal afferent is equal to the cupular deflection plus some noise. The SD of this Gaussian distributed noise on the canal afferent is one free parameter of their model. The second term says that the otolith afferent depends on head orientation and acceleration. This dependency implies an internal model of gravitoinertial force like that expressed in *Eq. 3*.

They calculate the particle filter prior distribution as follows

$$p(x_t|x_{t-1}) = p(\Omega)p(A)p(C_t|C_{t-1}, \Omega_t, \Omega_{t-1}) \qquad (13)$$

The first and second terms on the right-hand side are prior distributions for head angular velocity and linear acceleration, respectively. The authors define these to be zero-mean Gaussian distributions, thus suggesting that the most likely state of the head is zero or low angular velocity and linear acceleration. These priors are defined to be stationary distributions; this means that they do not change over time. Their SDs constitute the remaining two of three total free parameters of the model. The rightmost term in *Eq. 13* says that the current state of the cupula depends on the previous state of the cupula and on *the rate of change* of the angular velocity estimate. This distribution can therefore be calculated using an internal model of sensor dynamics, i.e., the canal transfer function.

In summary, the model has only three free parameters: noise on the canal signal (σ_V), variability of the angular velocity prior (σ_Ω), and variability of the linear acceleration prior (σ_A). With these three parameters fixed, model simulations generate results consistent with a variety of perceptual phenomena

including velocity storage, the somatogravic illusion, and illusory self-motion percepts arising from off-vertical axis rotation (Laurens and Droulez 2007).

Particle and Kalman filter models compared

Particle and Kalman filter models are very similar. Both incorporate internal models and both derive their motivation from analysis at the level of computational theory; they are probabilistic models that aim to generate the most likely or most probable state estimates and therefore they take into account the variability or noise associated with different sources of information. In fact, for linear system dynamics and Gaussian noise distributions, it can be shown that particle and Kalman filters generate virtually identical dynamic state estimates (J. Laurens, personal communication). Because they are probabilistic models, both the particle and Kalman filter models predict that the time constant for velocity storage should depend on the noise associated with the canal afferent. This is because noise on the canal signal is integrated over time, thus for a constant velocity rotation, the estimate of angular velocity will become less and less certain. As this happens, the estimate of angular velocity will converge toward zero, although different mechanisms drive this convergence for the two models (see below).

The advantage of the particle filter is that it can also handle non-Gaussian noise distributions and nonlinear system dynamics. These are significant limitations of the Kalman filter for modeling spatial orientation because the physical laws governing self-motion and orientation (*Eq. 4*) are nonlinear and Gaussian noise distributions are not guaranteed. Another difference is that the particle filter does not use explicit feedback circuits; it is more distributed. This distributed architecture, while computationally intensive, may be more suited to implementation by neural networks of the brain. We return to this topic in later sections.

Laurens and Droulez shared their implementation of the particle filter model with us and we used it to simulate velocity storage while varying the values of canal noise (σ_V) and the variability of the angular velocity prior (σ_Ω). The outcome of these simulations are shown in Fig. 6. As illustrated for the Kalman filter model in Fig. 5, the time constant depends on the noise associated with the canal signal, but it also depends on the other free parameter. In the Kalman filter model, this parameter is β, the bandwidth cutoff of the stimulus internal model (*Eq. 11*), whereas in the particle filter model this parameter is σ_W, the variability of the prior distribution for zero angular velocity.

These two parameters are functionally similar in that they serve to make the angular velocity estimate converge to zero in response to a constant angular velocity stimulus. They have been included because the goal is to characterize human behavior, which is quantitatively suboptimal for constant velocity rotation. This characteristic may be evidence of a compromise wherein accuracy of responses to uncommon (i.e., unnatural) stimuli, like constant velocity rotation, is sacrificed to maintain calibration of the system under more common, naturalistic conditions. In other words, when the canals are not being stimulated it is more likely that the head is stationary than rotating at constant velocity.

This rationale is the conceptual motivation of the Bayesian prior distribution for zero angular velocity included in the Laurens and Droulez (2007) particle filter model. However, the conceptual motivation for the stimulus internal model of the Kalman filter proposed by Borah et al. (1988) is somewhat different. It proposes that the nervous system maintains expectations about the interrelationship between angular velocity, acceleration, and jerk that depend on the
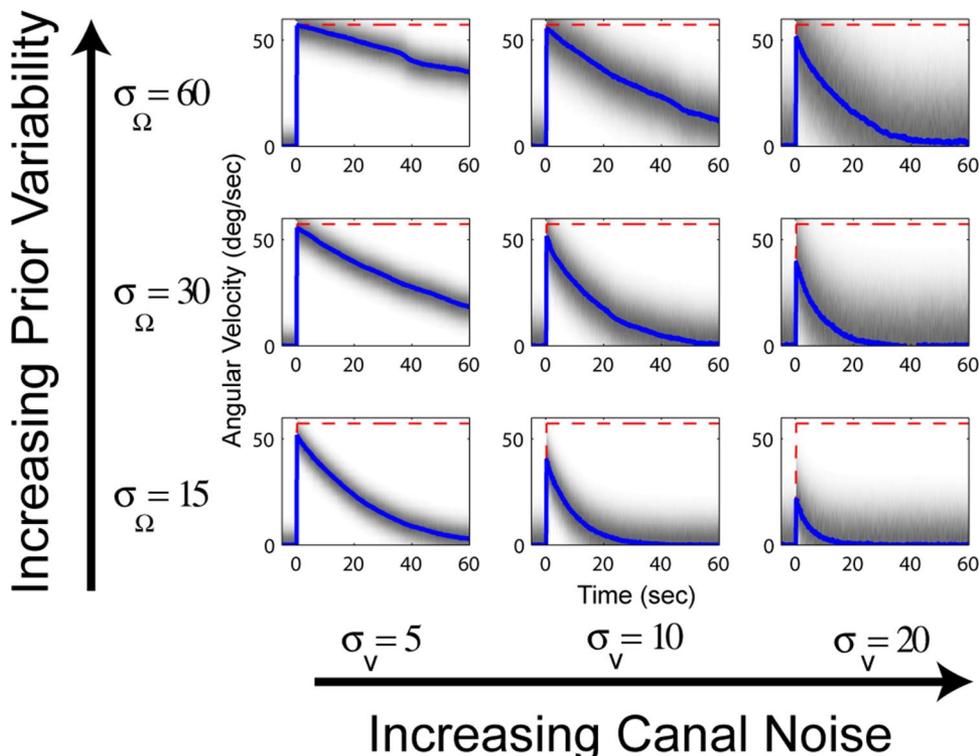

FIG. 6. In a particle filter model, the time constant of the angular velocity estimate (velocity storage) also depends on 2 parameters: the noise in the canal signal (σ_V) and the variability of the angular velocity prior (σ_Ω). The red dashed line illustrates the actual angular velocity and the blue lines show the time course of the angular velocity estimate. Values of canal noise (σ_V) increase on the columns from left to right from 5 to 20°/s. Values for prior variability (σ_Ω) increase on the rows from bottom to top from 15 to 60°/s. The time constant for velocity storage depends (approximately) on the ratio between σ_V and σ_Ω. The grayscale represents the variability associated with the estimate.

frequency of stimulation (*Eq. 11*). In retrospect, it could be that the stimulus internal model was introduced in the same spirit as the Bayesian prior. Higher β values correspond to a faster decay of the impulse response, which results in a lower amplitude angular velocity estimate (see Fig. 5). However, the construct is less intuitive than the idea of a Bayesian prior and seems less likely to accurately represent the computational goals of the system.

Interestingly, there is a parallel comparison to be drawn in the context of processing linear acceleration signals from the otoliths. Laurens and Droulez (2007) and MacNeilage et al. (2007) propose a Bayesian prior for zero or low inertial acceleration to account for the observation that humans tend to interpret constant linear acceleration as due to tilt (i.e., somatogravic illusion) because constant accelerating displacement is unlikely. An alternative explanation for this phenomenon is known as the frequency segregation model, which proposes that low- and high-frequency stimulation are interpreted as tilt and translation, respectively (Glasauer 1995; Mayne 1974; Seidman and Page 1996; Seidman et al. 1998). This can be thought of as an expectation on the frequency of otolith stimulation. Both models aim to explain the same observation, and given the right choice of parameters, both models will generate results consistent with human behavior. However, the question remains, which model more accurately represents the computation goals of the system? Does the nervous system maintain expectations about the amplitude of inertial acceleration, as suggested by the Bayesian prior, or the frequency of otolith stimulation, as suggested by frequency segregation?

In summary, the model proposed by Laurens and Droulez (2007) suggests that human vestibular processing for spatial orientation is governed by dynamic Bayesian inference and the model is implemented using the computationally flexible method of particle filtering. It incorporates internal models of sensor dynamics and physical laws as described in the preceding text. Noise associated with sensory signals is incorporated, and prior knowledge or experience is represented in the form of prior distributions. Similar spatial orientation priors for body orientation (de Vrijer et al. 2008; Eggert 1998; MacNeilage et al. 2007), linear acceleration (MacNeilage et al. 2007), and angular velocity (Paulin 2005) have been proposed by other authors. The model is dynamic and can generate perceptual estimates of position, velocity, and acceleration that evolve over time as new information is obtained and old information fades away; this is a distinct advantage over static Bayesian models. Yet the model could be made more realistic by adding noise on the otolith signal, and it could be made more comprehensive with the addition of other sensory information. In fact, a more recent modification of the model includes visual estimates of head motion in space (Laurens and Droulez 2008).

An important feature of any model is that it generates predictions that may be tested empirically, and the particle filter model does this. As noted in the preceding text, the dual dependency of the velocity storage time constant on canal noise and prior variability means that infinite combinations of these two parameters can yield the same time constant; for example, the time constants for the parameter combinations illustrated on one of the diagonals of Fig. 6 (lower left to top right) are approximately equal. However, note that the variability of the resulting angular velocity estimates (illustrated by the grayscale) are quite different, ranging from more precise in the lower left to less precise in the upper right. By empirically measuring the variability of the system's angular velocity estimate, it would be possible to calculate what the values of the canal noise and prior variability should be for this model. More generally, measuring variability or noise in computational systems, whether at the level of individual neurons or behavioral responses, is essential for developing, testing, and evaluating probabilistic theories of cognition (Faisal et al. 2008).

Another relevant feature of the particle filter model is its distributed nature, which contrasts with the explicit feedback architecture employed by the Kalman filter. With the particle filter, there is no error signal and no gain calculated. Interactions of different types of sensory information and estimates over time arise as a consequence of joint probability distributions and the recalculation of these distributions at each time step. This distributed quality mimics the computational architecture and function of actual neural networks in the brain. Particle filter implementations are computationally very intensive. Each probability distribution must be sampled many times at each time step to generate the distributions at the next time step with the number of samples depending on the number of particles used.

Notably, Kalman and particle filter approaches operate primarily at the level of computational theory. They seek statistically optimal state estimates given diverse, stochastic, and dynamic streams of sensory information. The computational hierarchy proposed by Marr (1982) emphasizes that such computational goals may be achieved using a number of alternative algorithms, and each of these algorithms may be physically implemented in a variety of ways. If the nervous system seeks statistically optimal solutions, it must approximate the function of a particle filter using representations and algorithms that may be realistically executed in real time by networks of neurons. Next we briefly review a few recent studies proposing neural implementations of Kalman and particle filters.

Neural implementations of Kalman and particle filters

Kalman and particle filters are manifestations of dynamic Bayesian inference. To perform Bayesian inference, the nervous system must have a way to represent probability distributions and a way to combine or multiply these distributions, as expressed in *Eq. 5*. Populations of sensory neurons tuned to a range of possible stimulus values are well-suited to representing full probability distributions (Foldiak 1993). It is generally believed that in the cortex, the response of each neuron to the same stimulus varies from presentation to presentation with Poisson-like statistics (Tolhurst et al. 1983). Ma, Beck, Latham, and Pouget (2006) point out that for probabilistic population codes composed of neurons with Poisson-like trial-to-trail variability, the peak of activity indicates the mean of the probability distribution and the gain of peak activity is inversely proportional to the variability of the corresponding Gaussian probability distribution. These authors demonstrate that the primary advantage of coding probability distributions in this way is that it makes combining (i.e., multiplying) probability distributions extremely straightforward; the population codes can simply be summed together. The linear combination of two or more population codes gives rise to com-

bined activity the gain of which is equal to the sum of the gains of the contributing population codes. Because the gain of the combined population code is greater, the variability of the corresponding Gaussian distribution is reduced, and the reduction is exactly that predicted by Bayesian integration.

This basic method of representing and combining Gaussian probability distributions with population codes composed of neurons with Poisson-like trial-to-trial variability has recently been extended to implement a Kalman filter (Deneve et al. 2007). In this model, several such population codes connect reciprocally with an intermediate network layer known as a *basis function layer*. Weights are applied to these reciprocal connections to ensure that, despite initially noisy inputs, activity patterns will converge to stable states; this architecture is known as an *attractor network*. The authors point out that this kind of neural network architecture is plausible because it is consistent with the noise statistics and neural interconnectivity observed in the cortex. Several predictions of this model are supported by experimental findings, including the observation that sensorimotor transformations are not necessarily computed using discrete reference frames (Avillac et al. 2005; Fetsch et al. 2007; Mullette-Gillman et al. 2005).

Deneve et al. (2007) conclude that the next step toward a more general and comprehensive neural theory of Bayesian inference requires a mechanism to encode arbitrary (non-Gaussian) probability distributions. Recall that dynamic Bayesian inference with non-Gaussian distributions can be achieved with a particle filter. So how could the nervous system implement a particle filter with networks of neurons? Paulin (2005) describes a neural architecture for which distributions need not be Gaussian. Specifically, Paulin (2005) claims that neural spikes (like individual particles in a particle filter) may be collected to reconstruct a Monte Carlo approximation to any probability density function. Furthermore he suggests a general mechanism for multiplying probability distributions; the postsynaptic firing thresholds can be adjusted such that the postsynaptic neuron only fires if two or more presynaptic neurons, each representing the probability of a specific state, are active at the same time. Given these methods for representing and multiplying probability distributions with spiking neurons, it is possible to implement a particle filter for dynamic Bayesian state estimation. Paulin (2005) eagerly points out that this technique can be applied to virtually any kind of state estimation problem and suggests that the neural architecture of the cerebellum is particularly well-suited to implementing this kind of neural computation.

Conclusion

Researchers have applied different computational approaches to the problem of estimating spatial orientation, and here we have tried to put these various contributions in context. We have attempted to classify each approach by describing the levels of analysis at which it operates. To demonstrate the conceptual motivations and mechanics of several approaches, we considered their application to modeling velocity storage. The Kalman and particle filter models suggest that velocity storage, and spatial orientation dynamics more generally, result from a complex web of constraints. Dynamic sensory signals must be integrated to generate useful estimates, but the integration process accumulates noise. When faced with uncer-

tainty, the system can rely on assumptions based on natural statistics and prior experience. Thus the dynamics of velocity storage and spatial orientation more generally have been shaped, through evolution, development, and experience, by natural statistics and computational compromises.

This conclusion emerged from considerations at the level of computational theory and does not specify the algorithm or implementation. For example, both Kalman and particle filters are models for dynamic Bayesian inference but they follow different algorithms; the Kalman filter is amenable to an analytic solution while the particle filter is a numerical simulation technique. However, if all noises are Gaussian and all system dynamics are linear, Kalman and particle filters yield identical state estimates. And just as the algorithm is independent of the computational goal, so is the implementation independent of the algorithm. For example, Kalman and particle filters can be implemented using Matlab or a variety of neural network architectures (Deneve et al. 2007; Paulin 2005). It is compelling to speculate how behavior resembling that of Kalman and particle filters can be implemented with networks of neurons. However, it will be very challenging to characterize and identify these processes at the neuronal level. Unlike previous neurophysiological research motivated by the transfer function approach, the neural substrates for dynamic Bayesian processing cannot be identified by simply measuring the mean activities of individual neurons. Instead to demonstrate how the brain accomplishes statistically optimal inference, it will be necessary to identify how entire neural populations represent sensory and prior information probabilistically. This intriguing field is just emerging and will likely be a target of intense research in the future.

APPENDIX A: TRANSFER FUNCTION TO STATE SPACE CONVERSION

Transfer function representation

The transfer function is obtained by Laplace transform of the linear ordinary differential equation (ODE) describing the underlying dynamics of the system. The transfer function simply connects the input (e.g., angular velocity) to the output (e.g., afferent response) without regards to other dynamic and internal states of the system. It is given by a proper rational polynomial

$$T(s) = \frac{y(s)}{u(s)} = \frac{\sum_{i=0}^{r-1} b_i s^i}{s^n + \sum_{i=0}^{n-1} c_i s^i}, \quad (r < n) \qquad (A1)$$

where $y(s)$ is the Laplace transform of the output and $u(s)$ that of the input to the system. The roots of the numerator polynomial are called the *zeros* of the system and that of the denominator are called the *poles* of the system. Together the zeros and poles describe the intensity of the dynamics, response times (time constants) and output of the system.

State space representation

In the state space representation, the current state of the system is *completely* described by a state vector which is an element of a vector space R^n where n is the order of the ODE. The elements of the state vector $X = \{x_1, x_2, \ldots, x_n\}^T$ describe various internal states of the

system. Hence the system can be thought of as residing in the vector space also called the *state space*. For example, a spring mass system the dynamics of which is described by a second-order ODE in terms of the external force is completely described by the knowledge of the position and displacement alone. All the higher derivatives including the acceleration can be expressed in terms of these variables at any time instant t. Hence in general, the number of independent variables required to completely represent a system whose dynamics is described by an nth order ODE is n. It is now possible to rewrite the dynamics of the system as a first order vector ODE quantifying the time rate of change of the state vector at every instant

$$\dot{X}(t) = AX(t) + bu(t)$$

$$y(t) = CX(t) \quad (A2)$$

where $y(t)$ and $u(t)$ are the values of the output and input at time t. Matrix A describes the main dynamics of the system (also called the *system matrix*); its eigenvalues describe the poles of the system. In *Eq. A2*, b describes the input dependence of the rate of change of the state, and C the state dependence of the output. The matrix A is $n \times n$, b is $n \times 1$, and C is $1 \times n$ for a single-input-single-output system. The conversion from transfer function to state space for a single-input-single-output system begins by identifying the coefficients of numerator and denominator polynomials (e.g., A1). The corresponding matrices A–C are then given by

$$A = \begin{bmatrix} 0 & 1 & 0 & \ldots & 0 \\ 0 & 0 & 1 & \ldots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & \ldots & 1 \\ -c_0 & -c_1 & -c_2 & \ldots & -c_{n-1} \end{bmatrix}$$

$$b = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 0 \\ 1 \end{bmatrix}$$

$$C = \begin{bmatrix} b_0 & b_1 & \ldots & b_{r-1} & 0 & \ldots & 0 \end{bmatrix} \quad (A3)$$

Improper transfer functions

An improper rational function is one in which the degree of numerator is equal to or greater than that of the denominator. For example

$$\frac{y(s)}{\omega(s)} = \frac{s}{s + T_c} \quad (A4)$$

represents such an improper transfer function for the semicircular canals, where ω is angular velocity and y is the afferent response. The state space conversion of this transfer function requires renaming the internal system states, making them different from the states of the proper rational system. For instance, the internal state would be the time integral of the afferent response not the afferent response itself. Such transfer functions are not conventionally used in systems theory because they

introduce additional terms in the state space model that void the input $\rightarrow$ state $\rightarrow$ output relation. The state space equations for improper rational functions have the form

$$\dot{X}(t) = AX(t) + bu(t)$$

$$y(t) = CX(t) + Du(t) \quad (A5)$$

The additional term D indicates that the output (e.g., afferent response) is directly affected by both the internal state ($\dot{X}(t)$) as well as the input (e.g., angular velocity).

APPENDIX B: IMPLEMENTATION OF KALMAN FILTER

In this section, we consider implementation of Kalman filters for discrete time systems and the generalization to continuous time case. For a detailed description of Kalman filter methods, see Maybeck (1979). Central to the implementation of the filter is the knowledge of the model parameters and the dynamics of the system. The system matrices A_k, B_k, and C_k define the input output relation, the input sequence $\{u_1, \ldots, u_k\}$ and the output map $\{y_1, \ldots, y_k\}$. The statistically optimal gain matrix L_k is computed as

$$L_k = A_k \Sigma_k C_k' [C_k \Sigma_k C_k' + \Theta_k]^{-1} \quad (B1)$$

where Θ_k is the covariance of the output noise and Σ_k is the covariance of the estimation error, also called the *error covariance*. It is defined as

$$\Sigma_k = \text{cov}(X_k - \hat{X}_k, \quad X_k - \hat{X}_k) \quad (B2)$$

The error covariance is calculated from the Riccati equation

$$\Sigma_{k+1} = A_k [\Sigma_k - \Sigma_k C_k' [C_k \Sigma_k C_k' + \Theta_k]^{-1} C_k \Sigma_k] A_k' + d_k \Xi_k d_k' \quad (B3)$$

Continuous time systems

The preceding implementation of Kalman filters for discrete time systems can be extended to continuous time systems (Kalman-Bucy) by suitable modifications to the error covariance calculation by the Riccati equation. The continuous time equation reads

$$\dot{\Sigma} = A(t)\Sigma(t) + \Sigma(t)A(t) + \Xi(t)d'(t)$$
$$- \Sigma(t)C'(t)\Theta^{-1}(t)C(t)\Sigma(t) \quad (B4)$$

In the implementation of the filter (Fig. B1), the delay element reduces to zero and the internal model is now a continuous time state space model. The corresponding continuous-time Kalman gain (analogous to *Eq. B1*) is now given by

$$L(t) = \Sigma(t)C(t)\Theta^{-1}(t) \quad (B5)$$

APPENDIX C: IMPLEMENTATION OF PARTICLE FILTER

Particle filters use a recursive algorithm that requires updating the previous state estimate to obtain the current state estimate. This is
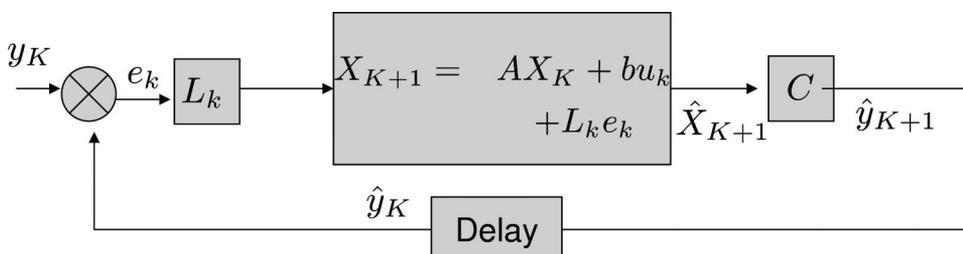

FIG. B1. The Kalman filter is implemented by the internal model parameters A–C and the computed Kalman gain L_k.

similar to the Kalman filter logic expressed in *Eq. 10*. The *previous estimate* depends on all previous information; this relationship is captured by the probability distribution $p(x_{t-1}|y_{1:t-1})$. The *current estimate* depends on all previous and current information and can be expressed by the probability distribution $p(x_t|y_{1:t})$. For a detailed description of the particle filter method, see Maskell and Gordon (2002).

The particle filter algorithm is usually expressed in terms of a prediction step and an updating step. The *prediction step* expresses the probability of the current state (x_t) given the previous information ($y_{1:t-1}$)

$$p(x_t|y_{1:t-1}) = \int p(x_{t-1}|y_{1:t-1})p(x_t|x_{t-1})dx_{t-1} \qquad (C1)$$

The first term on the right-hand side of this equation is the previous estimate that the particle filter aims to update. The second term, $p(x_t|x_{t-1})$, says that the current state depends on the previous state and therefore contains information about both the *prior* and the rate of change of the state. The probability distribution for the prediction (the term on the left) is obtained by integrating over all possible values of the previous state (x_{t-1}).

The *updating step* expresses the probability of the *current estimate*, $p(x_t|y_{1:t})$, in terms of the prediction, $p(x_t|y_{1:t-1})$

$$p(x_t|y_{1:t}) = \frac{p(y_t|x_t)p(x_t|y_{t-1})}{\int p(y_t|x_t)p(x_t|y_{1:t-1})dx_t} \qquad (C2)$$

The additional term in this equation, $p(y_t|x_t)$, is the probability of the current information given the current state, the particle filter version of the likelihood function. It depends on knowledge (i.e., an internal model) of the sensor. Note that the current estimate can be thought of as the posterior probability distribution in the Bayesian framework because the updating step makes use of Bayes' rule to express the current estimate in terms of the other two distributions. The current estimate can be expressed as $p(x_t|y_{1:t}) = p[(x_t|y_{1:t-1})|y_t]$. Applying Bayes' rule (*Eq. 5*) to the term on the right we obtain $p(x_t|y_{1:t}) = \{p[y_t|(x_t|y_{1:t-1})]p(x_t|y_{1:t-1})\}/[p(y_t)]$, which can be rewritten as *Eq. C2*.

As described in the text, particle filtering is a numerical simulation technique in which all of the preceding probabilities are calculated based on samples of these same distributions from the previous time step. The simulation steps involved in implementing the particle filter are as follows. *1*) The prior is sampled into N particles with corresponding weights (probabilities) w^i. The particles can now be thought of as possible values of the state estimate at time $t = 0$, e.g., an initial velocity estimate. *2*) Knowing the estimate at time t and the system dynamics expressed in *Eqs. 12* and *13,* the particles can be propagated forward to obtain the estimate at time $t + 1$. *3*) At this point, it is likely that the new probability distribution is not captured accurately. For example, it is possible for the values (estimates) that are less probable to be represented more than those with greater probability. Hence a resampling is carried out where the probabilities are normalized. The estimates that are more probable are now encoded in more particles than those that are less probable.

The renormalization step ensures that improbable estimates are systematically eliminated, thus ensuring the accuracy and reliability of the computed probability distribution. *Steps 1–3* are iterated the required number of times to obtain the probability distribution at the end of the *k*th time step. The fidelity of any particle filter implementation depends on the number of particles used and the duration of each time step. More particles and shorter time steps will more accurately reproduce the dynamics of the system being modeled but will demand more computational resources.

R E F E R E N C E S

Angelaki DE. Eyes on target: what neurons must do for the vestibuloocular reflex during linear motion. *J Neurophysiol* 92: 20–35, 2004.
Angelaki DE, Cullen KE. Vestibular system: the many facets of a multimodal sense. *Annu Rev Neurosci* 31: 125–150, 2008.
Angelaki DE, Dickman JD. Spatiotemporal processing of linear acceleration: primary afferent and central vestibular neuron responses. *J Neurophysiol* 84: 2113–2132, 2000.
Angelaki DE, McHenry MQ, Dickman JD, Newlands SD, Hess BJ. Computation of inertial motion: neural strategies to resolve ambiguous otolith information. *J Neurosci* 19: 316–327, 1999.
Angelaki DE, Shaikh AG, Green AM, Dickman JD. Neurons compute internal models of the physical laws of motion. *Nature* 430: 560–564, 2004.
Avillac M, Denève S, Olivier E, Pouget A, Duhamel JR. Reference frames for representing visual and tactile locations in parietal cortex. *Nat Neurosci* 8: 941–949, 2005.
Awalt B, Turner D, Miller R, Chen W-Z. Extended Kalman filter applications to multi-vehicle UAV cooperative controls. *Second AIAA Unmanned Unlimited Systems, Technologies, and Operations,* San Diago, CA, 2003.
Berger JO. *Statistical Decision Theory and Bayesian Analysis* (2nd ed.). New York: Springer Verlag, 1985.
Borah J, Young, LR, Curry RE. Optimal estimator model for human spatial orientation. *Ann NY Acad Sci* 545: 51–73, 1988.
Bortolami SB, Rocca S, Daros S, DiZio P, Lackner JR. Mechanisms of human static spatial orientation. *Exp Brain Res* 173: 374–388, 2006.
Bos JE, Bles W. Modelling motion sickness and subjective vertical mismatch detailed for vertical motions. *Brain Res Bull* 47: 537–542, 1998.
Bos JE, Bles W. Theoretical considerations on canal-otolith interaction and an observer model. *Biol Cybern* 86: 191–207, 2002.
Cohen B, Matsuo V, Raphan T. Quantitative analysis of the velocity characteristics of optokinetic nystagmus and optokinetic after-nystagmus. *J Physiol* 270: 321–344, 1977.
Cullen KE, Roy JE. Signal processing in the vestibular system during active versus passive head movements. *J Neurophysiol* 91: 1919–1933, 2004.
de Vrijer M, Medendorp WP, Van Gisbergen JA. Shared computational mechanism for tilt compensation accounts for biased verticality percepts in motion and pattern vision. *J Neurophysiol* 99: 915–930, 2008.
Denève S, Duhamel J-R, Pouget A. Optimal sensorimotor integration in recurrent cortical networks: a neural implementation of Kalman filters. *J Neurosci* 27: 5744–5756, 2007.
Dickman JD, Angelaki DE. Vestibular convergence patterns in vestibular nuclei neurons of alert primates. *J Neurophysiol* 88: 3518–3533, 2002.
Eggert T. *Der Einfluss orientierter Texturen auf die subjektive visuelle Vertikale und seine systemtheoretische Analyse* (PhD thesis). Munich, Germany: Munich Technical Univ., 1998.
Ernst MO, Banks MS, Humans integrate visual and haptic information in a statistically optimal fashion, *Nature* 415: 429–433, 2002.
Faisal AA, Selen LP, Wolpert DM. Noise in the nervous system. *Nat Rev Neurosci* 9: 292–303, 2008.
Fetsch CR, Wang S, Gu Y, Deangelis GC, Angelaki DE. Spatial reference frames of visual, vestibular, and multimodal heading signals in the dorsal subdivision of the medial superior temporal area. *J Neurosci* 27: 700–712, 2007.
Fernández C, Goldberg JM. Physiology of peripheral neurons innervating otolith organs of the squirrel monkey. I. Response to static tilts and to long-duration centrifugal force. *J Neurophysiol* 39: 970–984, 1976.
Foldiak P. The "ideal homunculus": statistical inference from neural population responses. In: *Computation and Neural Systems,* edited by Eeckman F, Bower J. Norwell, MA: Kluwer Academic Publishers, 1993 p. 55–60.
Fukushima K. The interstitial nucleus of Cajal in the midbrain reticular formation and vertical eye movement. *Neurosci Res* 10: 159–187, 1991.

Galiana HL, Outerbridge JS. A bilateral model for central neural pathways in vestibuloocular reflex. *J Neurophysiol* 51: 210–241, 1984.

Gaskill JD. *Linear Systems, Fourier Transforms, and Optics.* New York: Wiley, 1978.

Ghahramani Z, Wolpert DM, Jordan MI. Computational models of sensorimotor integration. In: *Self-Organization, Computational Maps, and Motor Control,* edited by Morasso PG, Sanguineti V. Amsterdam: Elsevier, 1997, p. 117–147.

Ghasia FF, Meng H, Angelaki DE. Neural correlates of forward and inverse models for eye movements: evidence from three-dimensional kinematics. *J Neurosci* 28: 5082–5087, 2008.

Gibson JJ. *The Perception of the Visual World.* Boston, MA: Houghton Mifflin, 1950.

Glasauer S. Linear acceleration perception: frequency dependence of the hilltop illusion. *Acta Otolaryngol* 520, Suppl 1: 37–40, 1995.

Glasauer S. Current models of the ocular motor system. *Dev Ophthalmol* 40: 158–174, 2007.

Goldberg JM, Fernandez C. Physiology of peripheral neurons innervating semicircular canals of the squirrel monkey. I. Resting discharge and response to constant angular accelerations. *J Neurophysiol* 34: 635–660, 1971.

Goldman MS, Kaneko CR, Major G, Aksay E, Tank DW, Seung HS. Linear regression of eye velocity on eye position and head velocity suggests a common oculomotor neural integrator. *J Neurophysiol* 88: 659–665, 2002.

Green AM, Angelaki DE. Resolution of sensory ambiguities for gaze stabilization requires a second neural integrator. *J Neurosci* 23: 9265–9275, 2003.

Green AM, Angelaki DE. An integrative neural network for detecting inertial motion and head orientation. *J Neurophysiol* 92: 905–925, 2004.

Green AM, Galiana HL. Hypothesis for shared central processing of canal and otolith signals. *J Neurophysiol* 80: 2222–2228, 1998.

Green AM, Meng H, Angelaki DEA. reevaluation of the inverse dynamic model for eye movements. *J Neurosci* 27:1346–1355, 2007.

Green AM, Shaikh AG, Angelaki DE. Sensory vestibular contributions to constructing internal models of self-motion. *J Neural Eng* 2: S164–179, 2005.

Gu Y, Watkins PV, Angelaki DE, DeAngelis GC. Visual and nonvisual contributions to three-dimensional heading selectivity in the medial superior temporal area. *J Neurosci* 26: 73–85, 2006.

Hillis JM, Watt SJ, Landy MS, Banks MS. Slant from texture and disparity cues: optimal cue combination. *J Vis* 4: 967–992, 2004.

Kuo AD. An optimal control model for analyzing human postural balance. *IEEE Trans Biomed Eng* 42: 87–101, 1995.

Kuo AD. An optimal state estimation model of sensory integration in human postural balance. *J Neural Eng* 2: S235–S249, 2005.

Laurens J, Droulez J. Bayesian processing of vestibular information. *Biol Cybern* 96: 389–404, 2007.

Laurens J, Droulez J. Bayesian modeling of visuo-vestibular interactions. In: *Probabilistic Reasoning and Decision Making in Sensory-Motor Systems,* edited by Bessière P, Laugier C, Siegwart R. New York: Springer, 2008, p. 291–313.

Ma WJ, Beck JM, Latham PE, Pouget A. Bayesian inference with probabilistic population codes. *Nat Neurosci* 9: 1432–1843, 2006.

MacNeilage PR, Banks MS, Berger DR, Bülthoff HH. A Bayesian model of the disambiguation of gravitoinertial force by visual cues. *Exp Brain Res* 179: 263–290, 2007.

Marr D. *Vision:* A Computational Investigation into the Human Representation and Processing of Visual Information. New York: Freeman, 1982.

Maskell S, Gordon N. A tutorial on particle filters for on-line nonlinear/non-gaussian bayesian tracking. *IEEE Trans Signal Process* 50:174–188, 2002.

Maybeck PS. *Stochastic Models, Estimation and Control.* New York: Academic, 1979.

Mayne R. A systems concept of the vestibular organs. In: *Handbook of Sensory Physiology. Vestibular System. Psychophysics, Applied Aspects and General Interpretations,* edited by Kornhuber H. Berlin: Springer-Verlag, 1974, vol. VI, part 2, p. 493–580.

McFarland JL, Fuchs AF. Discharge patterns in nucleus prepositus hypoglossi and adjacent medial vestibular nucleus during horizontal eye movement in behaving macaques. *J Neurophysiol* 68: 319–332, 1992.

Meng H, May PJ, Dickman JD, Angelaki DE. Vestibular signals in primate thalamus: properties and origins. *J Neurosci* 27: 13590–13602, 2007.

Merfeld DM. Modeling the vestibulo-ocular reflex of the squirrel monkey during eccentric rotation and roll tilt. *Exp Brain Res* 106: 123–134, 1995.

Merfeld DM, Park S, Gianna-Poulin C, Black FO, Wood S. Vestibular perception and action employ qualitatively different mechanisms. I. Frequency response of VOR and perceptual responses during translation and tilt. *J Neurophysiol* 94: 186–198, 2005.

Merfeld DM, Young LR, Oman CM, Shelhamer MJ. A multidimensional model of the effect of gravity on the spatial orientation of the monkey. *J Vestib Res* 3: 141–161, 1993.

Merfeld DM, Zupan L, Peterka RJ. Humans use internal models to estimate gravity and linear acceleration. *Nature* 398: 615–618, 1999.

Mittelstaedt H. A new solution to the problem of the subjective vertical. *Naturwissenschaften* 70: 272–281, 1983.

Mullette-Gillman OA, Cohen YE, Groh JM. Eye-centered, head-centered, and complex coding of visual and auditory targets in the intraparietal sulcus. *J Neurophysiol* 94: 2331–2352, 2005.

Oman CM. A heuristic mathematical model for the dynamics of sensory conflict and motion sickness. *Acta Otolaryngol Suppl* 392: 1–44, 1982.

Paulin MG. Evolution of the cerebellum as a neuronal machine for Bayesian state estimation. *J Neural Eng* 2: S219–234, 2005.

Raphan T, Cohen B. The vestibulo-ocular reflex in three dimensions. *Exp Brain Res* 145: 1–27, 2002.

Raphan T, Matsuo V, Cohen B. A velocity storage mechanism responsible for optokinetic nystagmus (OKN), optokinetic after-nystagmus (OKAN) and vestibular nystagmus. In: *Control of Gaze by Brain Stem Neurons, Developments in Neuroscience,* edited by Baker R, Berthoz A. Amsterdam: Elsevier/North-Holland Biomedical, 1977, vol. 1, p. 49–58.

Raphan T, Matsuo V, Cohen B. Velocity storage in the vestibulo-ocular reflex arc (VOR). *Exp Brain Res* 35: 229–248, 1979.

Robinson DA. Vestibular and optokinetic symbiosis: an example of explaining by modeling. In: *Control of Gaze by Brain Stem Neurons, Developments in Neuroscience,* edited by Baker R, Berthoz A. Amsterdam: Elsevier/North-Holland Biomedical, 1977, vol. 1, p. 49–58.

Robinson DA. The use of control systems analysis in the neurophysiology of eye movements. *Annu Rev Neurosci* 4: 463–503, 1981.

Robinson FR, Fuchs AF. The role of the cerebellum in voluntary eye movements. *Annu Rev Neurosci* 24: 981–1004, 2001.

Scudder CA, Fuchs AF. Physiological and behavioral identification of vestibular nucleus neurons mediating the horizontal vestibuloocular reflex in trained rhesus monkeys. *J Neurophysiol* 68: 244–264, 1992.

Seidman SH, Paige GD. Perception and eye movement during low-frequency centripetal acceleration. *Ann NY Acad Sci* 781: 693–695, 1996.

Seidman SH, Telford L, Paige GD. Tilt perception during dynamic linear acceleration. *Exp Brain Res* 119: 307–314, 1998.

Shaikh AG, Green AM, Ghasia FF, Newlands SD, Dickman JD, Angelaki DE. Sensory convergence solves a motion ambiguity problem. *Curr Biol* 15: 1657–1662, 2005.

Si X, Angelaki DE, Dickman JD. Response properties of pigeon otolith afferents to linear acceleration. *Exp Brain Res* 117: 242–250, 1997.

Skavenski AA, Robinson DA. Role of abducens neurons in vestibuloocular reflex. *J Neurophysiol* 36: 724–738, 1973.

Spingarn K. Passive position location estimation using the extended Kalman filter. *IEEE Trans Aerospace Electron Syst* AES-23: 558–567, 1987.

Tin C, Poon CS. Internal models in sensorimotor integration: perspectives from adaptive control theory. *J Neural Eng* 2: S147–163, 2005.

Todorov E. Optimality principles in sensorimotor control. *Nat Neurosci* 7: 907–915, 2004.

Tolhurst DJ, Movshon JA, Dean AF. The statistical reliability of signals in single neurons in cat and monkey visual cortex. *Vision Res* 23: 775–785, 1983.

van der Kooij H, Jacobs R, Koopman B, Grootenboer H. A multisensory integration model of human stance control. *Biol Cybern* 80: 299–308, 1999.

van der Kooij H, Jacobs R, Koopman B, van der Helm F. An adaptive model of sensory integration in a dynamic environment applied to human stance control. *Biol Cybern* 84: 103–115, 2001.

Waespe W, Büttner U, Henn V. Visual-vestibular interaction in the flocculus of the alert monkey. I. Input activity. *Exp Brain Res* 43: 337–348, 1981.

Waespe W, Henn V. Vestibular nuclei activity during optokinetic after-nystagmus (OKAN) in the alert monkey. *Exp Brain Res* 30: 323–330, 1977a.

Waespe W, Henn V. Neuronal activity in the vestibular nuclei of the alert monkey during vestibular and optokinetic stimulation. *Exp Brain Res* 27: 523–538, 1977b.

Waespe W, Henn V. The velocity response of vestibular nucleus neurons during vestibular, visual, and combined angular acceleration. *Exp Brain Res* 37: 337–347, 1979.

Waespe W, Henn V. Visual-vestibular interaction in the flocculus of the alert monkey. II. Purkinje cell activity. *Exp Brain Res* 43: 349–360, 1981.

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

Wolpert DM. Probabilistic models in human sensorimotor control. *Hum Mov Sci* 26: 511–524, 2007.

Wolpert DM, Ghahramani Z. Computational principles of movement neuroscience. *Nat Neurosci Suppl* 3: 1212–1217, 2000.

Wolpert DM, Kawato M. Multiple paired forward and inverse models for motor control. *Neural Netw* 11: 1317–1329, 1998.

Young LR. *On Visual-Vestibualr Interaction: Proc. Fifth Symposium on the Role of the Vestibular Organs in Space Exploration,* Houston, TX: NASA, 1970. (Spec. NASA Rep. SP-314, p. 205–210).

Yakusheva TA, Shaikh AG, Green AM, Blazquez PM, Dickman JD, Angelaki DE. Purkinje cells in posterior cerebellar vermis encode motion in an inertial reference frame. *Neuron* 54: 973–985, 2007.

Zupan LH, Merfeld DM. Human ocular torsion and perceived roll responses to linear acceleration. *J Vestib Res* 15: 173–183, 2005.

Zupan, LH, Merfeld, DM, Darlot, C. Using sensory weighting to model the influence of canal, otolith and visual cues on spatial orientation and eye movements. *Biol Cybern* 86: 209–230, 2002.