Biol. Cybern. 92, 147–163 (2005)
DOI 10.1007/s00422-004-0535-x

Biological Cybernetics

Comparing internal models of the dynamics of the visual environment

Sean Carver[1], **Tim Kiemel**[1,2], **Herman van der Kooij**[3], **John J. Jeka**[1,4]

[1] Department of Kinesiology, University of Maryland, College Park, MD 20742, USA
[2] Department of Biology, University of Maryland, College Park, MD, USA
[3] Institute for Biomedical Technology, University of Twente, Enschede, The Netherlands
[4] Program in Neuroscience and Cognitive Science, University of Maryland, College Park, MD, USA

Received: 16 June 2004 / Accepted: 17 November 2004 / Published online: 9 February 2005

Abstract. It is well known that the human postural control system responds to motion of the visual scene, but the implicit assumptions it makes about the visual environment and what quantities, if any, it estimates about the visual environment are unknown. This study compares the behavior of four models of the human postural control system to experimental data. Three include internal models that estimate the state of the visual environment, implicitly assuming its dynamics to be that of a linear stochastic process (respectively, a random walk, a general first-order process, and a general second-order process). In each case, all of the coefficients that describe the process are estimated by an adaptive scheme based on maximum likelihood. The fourth model does not estimate the state of the visual environment. It adjusts sensory weights to minimize the mean square of the control signal without making any specific assumptions about the dynamic properties of the environmental motion.

We find that both having an internal model of the visual environment and its type make a significant difference in how the postural system responds to motion of the visual scene. Notably, the second-order process model outperforms the human postural system in its response to sinusoidal stimulation. Specifically, the second-order process model can correctly identify the frequency of the stimulus and completely compensate so that the motion of the visual scene has no effect on sway. In this case the postural control system extracts the same information from the visual modality as it does when the visual scene is stationary. The fourth model that does not simulate the motion of the visual environment is the only one that reproduces the experimentally observed result that, across different frequencies of sinusoidal stimulation, the gain with respect to the stimulus drops as the amplitude of the stimulus increases but the phase remains roughly constant. Our results suggest that the human postural control system does not estimate the state of the visual environment to respond to sinusoidal stimuli.

Correspondence to: S. Carver
(e-mail: sc350@umail.umd.edu,
Tel.: +1-301–4058272, Fax: +1-301–4055578)

1 Introduction

Significant progress in understanding the human postural control system over the last 10–20 years has emphasized the need to understand its inherently adaptive nature. It is now a generally held view that visual, vestibular, and somatosensory inputs are dynamically reweighted to maintain upright stance as environmental or nervous system conditions change (Horak and Macpherson 1996; Kiemel et al. 2002; Oie et al. 2002; Peterka 2002; Shumway-Cook and Woollacott 2001) Environmental changes such as moving from a light to a dark environment or from a fixed to a moving support surface (e.g., onto a moving walkway at an airport) require an updating of sensory weights to current conditions so that muscular commands are based on the most precise and reliable sensory information available (Teasdale et al. 1991; Wolfson et al. 1985; Woollacott et al. 1986). The emphasis on adaptive processes has highlighted the need for nonlinear models to account for multisensory reweighting. One approach to this problem has involved the implementation of an internal model (Van der Kooij et al. 2001).

An internal model simulates within the nervous system some natural process, be it internal or external to the body (Wolpert et al. 1995). An example of how such simulations can facilitate motor control is given by a Kalman filter. Mathematical models of posture that include state estimation typically implement some variation of this optimal estimator (Van der Kooij et al. 1999; Kiemel et al. 2002). A Kalman filter simulates the dynamics of the body and its sensors to predict its afferent sensory signals. The Kalman filter then compares its predictions to the actual signals to derive optimal state estimates given the statistics of the noise in the system. Thus an internal model is typically assumed to underlie the nervous system's state estimation needed to control posture.

State estimation for posture traditionally involves simulating the body and sensor dynamics. Less prominent is the idea that the nervous system simulates the dynamics of the environment, although it has been indirectly discussed in the posture literature for some time. Early studies (Brandt et al. 1973, 1975)

proposed that the perceptual system underlying postural control distinguishes objects in the environment (foreground) from the overall environmental motion (background) based upon evidence of visually induced self-motion perception (i.e., vection). For instance, vection is dependent upon the number of moving contrasts (stripes or dots) within a rotating visual display (Brandt et al. 1975). With very few moving contrasts in the visual display, vection is not reported, indicating that self-motion is distinguishable from environmental motion. However, vection is experienced more reliably as the number of moving contrasts increases. Similar effects have been observed by controlling the area (i.e., small vs. large) of the moving contrasts (Brandt et al. 1973). More recent experimental results from Ravaioli et al. (2004) also suggest that the nervous system decomposes relative motion between the visual scene and body sway into environmental and self-motion. Standing subjects were presented with visual displays whose motion consisted of both an oscillation and a constant velocity translation that varied between conditions. Gain relative to the oscillation did not initially decrease as the translation velocity increased but decreased significantly at higher translation speeds. These results suggest that the postural control system can selectively compensate for translation at low speeds. The authors proposed that an internal model, having an estimate of the visual scene translation velocity, simulates the translatory motion of the environment. Their proposed neural controller subtracts the simulated motion from the visual measurement to give an estimate of body motion.

Van der Kooij et al. (2001) was the first to propose an internal model of the environment as part of the postural control system by expanding the state estimation problem to include estimation of properties of the environment. Their internal model derives its prediction of the sensory signals not only by simulating the dynamics of the body and its sensors, but also by simulating the dynamics of the environment. Specifically, the internal model of their postural control system simulates the motion of the visual scene, the translation and rotation of the support base, and the dynamics of an externally applied force.

To implement a simulation of the environment, the nervous system must have a model of how it changes. The modeled nervous system of Van der Kooij et al. (2001) implicitly assumes the dynamics of the position of the visual scene (in our notation x_v) to be a random walk. This gives assumed (but possibly incorrect) state equations for x_v that can be included with the Kalman filter's equations for the dynamics of the body and sensors. Thus the nervous system keeps track of, utilizes, and updates an estimate of x_v, called $\hat{x}_v$. With this estimate the internal model predicts the sensory measurements. Based on the properties of a random walk, the internal model considers motion of x_v in either direction to be equally likely. Therefore, it changes $\hat{x}_v$ based on discrepancies between the actual sensory measurements and its predictions.

The update to $\hat{x}_v$ will depend not only upon the discrepancies, but also upon assumed noise statistics of the system, including an estimate of the noise level associated with the visual scene random walk dynamics (in our notation d_1). This estimate is a second quantity associated with the visual environment that the modeled nervous system maintains and updates. Van der Kooij et al. (2001) chose the adaptive scheme of Myers and Tapley (1976) to estimate d_1. It should be noted that if a random walk correctly models the visual dynamics, then the nervous system's estimate of d_1 will approximate its true value. Once this happens, the internal model of the environment will match reality and the state estimates upon which the control is based will be optimal in the sense made precise by optimal control theory (Bryson and Ho 1975).

On the other hand, if the environmental dynamics do not match a random walk, the nervous system's estimate of d_1 will depend upon the (possibly changing) properties of the visual stimulus. Thus the nervous system's behavior will adapt to changing environmental conditions. But does the model adapt in the same way as the human postural control system?

In many respects the model of Van der Kooij et al. (2001) succeeded in reproducing the behavior of the human postural control system. Notably, the authors found good quantitative agreement with the experimental data of Peterka and Benolken (1995) with regard to gain in response to a sinusoidal stimulus across different frequencies and amplitudes of the stimulus. Qualitatively the model predicts that gain drops as a function of increasing stimulus amplitude across different stimulus frequencies – correctly reproducing this well-known experimental result (Peterka and Benolken 1995). Van der Kooij et al. (2001) found their model agreed in other respects with the data of Peterka and Benolken (1995). Among other things they found that without the vestibular sensors, their model reproduced reasonably well gain data from vestibular loss patients.

This paper tests more rigorously the hypothesis that an internal model of the environment underlies the adaptation of the human postural control system to changing environmental conditions. We use observations not considered by Van der Kooij et al. (2001) to further constrain the structure and parameterization of the model. Specifically, we consider the phase of the model in response to a sinusoidal visual stimulus as a function of the frequency and amplitude of the stimulus. Indeed, system identification customarily involves both gain and phase. Numerous studies (Peterka and Benolken 1995; Oie et al. 2002) have found that phase for human subjects in response to a sinusoidal stimulus, across different frequencies, is roughly constant as a function of amplitude of the stimulus. At the parameter values we tested we found that the model of Van der Kooij et al. (2001) did not reproduce the phase response seen with human subjects as a function of amplitude. Specifically, at the values of the parameters we tested, we found a substantial increase in phase (of about 150° with a 0.2-Hz stimulus between amplitudes of 0.2° and 5°).

To address the issue of proper phase behavior, we were led to consider simpler models than the ones studied in Van der Kooij et al. (2001). Because of their simplicity, our models were easier to understand and faster to simulate. With formulas that allowed us to approximate each model's transfer function in less than a second, we were

able to thoroughly explore each model's parameter space. We believe that different insights can come from studying simpler models than from studying more realistic ones and that both approaches should be taken together.

The models we present in this paper are simpler than the model of Van der Kooij et al. (2001) in four respects. First, the support base is fixed, and we do not include an external force: the only variation in the environment is the motion of the visual scene. Second, instead of a full set of five sensors, we implement only two: a visual sensor and a nonvisual sensor. Third, we do not model sensory dynamics. The transfer function of each sensor (with respect to velocity) is 1. Fourth, there is no time delay in our model. As described below, our models also differ in several other respects from the model of Van der Kooij et al. (2001).

Because we have simplified the model substantially, we do not expect the model to agree quantitatively with experimental data. Instead, we test for qualitative agreement. Specifically, we seek to reproduce the experimental result that across different frequencies of sinusoidal stimulation the gain with respect to the stimulus drops as the amplitude of the stimulus increases, but the phase remains roughly constant. This constraint proved to be difficult to meet.

After considering the behavior of our simple postural model with a random walk internal model of the environment, we compare this behavior to its behavior with two more sophisticated internal models. While these more sophisticated internal models do not produce qualitative agreement with experimental results, their behavior yields interesting insights into how a postural control system is theoretically capable of performing in response to sinusoidal stimulation.

Finally, we compare the behavior of the three postural models with internal models simulating the visual environment in order to estimate its state to the behavior of a postural model with no internal model simulating the environment. This nonsimulating and nonestimating scheme is the only one we found that produces qualitative agreement with experimental data. We have found that, with a stochastic process internal model simulating the visual environment to estimate its state, substantial drops in gain with increasing stimulus amplitude tend to come with substantial rises in phase at some stimulus frequencies. This tendency could be reduced by tuning the parameters, but we were unable to eliminate it.

2 Methods

2.1 Postural models

We modeled the standing human body as a linearized inverted pendulum whose stability was maintained by a proportional-derivative controller that depended upon an estimate of the body's position and velocity.

The modeled nervous system had two sensors. The first sensor was visual (it output the body's velocity relative to the visual scene plus noise); the second sensor was nonvisual (it output the body's true velocity plus noise). Thus neither sensor included dynamics: the transfer function of both, with respect to velocity, was 1. We chose velocity coupling for our sensors rather than position or acceleration coupling (or some combination of the three) as a simplification based on the hypothesis that velocity is the most accurate form of information for postural control (Jeka et al. 2004; Kiemel et al. 2002).

We considered three postural models that based their response to the motion of the visual scene on the estimates of progressively more sophisticated internal models of the visual environment. In addition, we considered one nonestimating postural model whose response to visual stimulation was not based on estimates of the properties of visual environment. We expressed each of the four systems as a stochastic differential equation.

Each model described both the computations of the neural controller as well as the dynamics of the body. In other words, each model's stochastic differential equation contained some components describing the body dynamics as well as some components describing the dynamics of the neural computations. An important constraint on a well-stated postural model is that the computations performed by the neural controller depend on the body dynamics and the noise in the system only through a function that defines the sensory measurement.

To make this statement precise, the state of the body was described by a vector x whose dynamics were given by the following stochastic differential equation in x, the ankle torque applied by the controller u, and the process noise ξ_p:

$$\dot{x}(t) = Fx(t) + Gu(t) + \xi_p(t)\,. \tag{1}$$

The components of x were the angular position and angular velocity of the body. The sensory measurement z was given in terms of x, the visual scene position x_v, and measurement noise ξ_m by the equation:

$$z(t) = Hx(t) + E\dot{x}_v(t) + \xi_m(t)\,. \tag{2}$$

The first component of z represented the visual measurement, the second component the nonvisual measurement. We assumed $\xi_p(t)$ and $\xi_m(t)$ were independent Gaussian white noise processes with spectral densities Q and R, respectively. For each of the four postural models we added components (a neural controller) to (1) that determined the dynamics of $u(t)$. To satisfy the constraint that the postural model be well stated, we only allowed the neural controller components to depend upon the quantities in (1) through the measurement $z(t)$. More precisely, the neural controller had access to both $z(t)$ and coefficients of the matrices F, G, Q, H, E, and R, but it did not have direct access to the processes $x(t)$, $\xi_p(t)$, $x_v(t)$, and $\xi_m(t)$. The matrices in (1) and (2) were given by

$$F = \begin{bmatrix} 0 & 1 \\ \gamma & 0 \end{bmatrix}, \quad G = \begin{bmatrix} 0 \\ 1 \end{bmatrix}, \quad Q = \begin{bmatrix} 0 & 0 \\ 0 & \sigma^2 \end{bmatrix}, \quad H = \begin{bmatrix} 0 & 1 \\ 0 & 1 \end{bmatrix},$$
$$E = \begin{bmatrix} -1 \\ 0 \end{bmatrix}, \quad R = \begin{bmatrix} \sigma_{21}^2 & 0 \\ 0 & \sigma_{22}^2 \end{bmatrix}, \tag{3}$$

where γ, σ, σ_{21}, and σ_{22} are parameters of the model.

The parameter γ is related to the torsional stiffness of the system. Specifically, γ is the torsional stiffness divided

by the moment of inertia, which is the proportionality factor that relates body angle to the angular acceleration caused by that angle. The parameter γ can be expressed as

$$\gamma = \frac{mgh}{J} - \gamma_p ,$$

where m is the mass of the body, g is the acceleration due to gravity, h is the height of the center of mass above the ankles, J is the moment of inertia of the body about the ankles, and γ_p is the passive component to γ. The first term of this expression can be calculated for each subject from anthropometric quantities that can be measured precisely. Values of this term used by postural models in the literature have included $9.83\,s^{-2}$ (Peterka 2000) and $7.6\,s^{-2}$ (Kiemel et al. 2002). On the other hand, there is no agreement about the value of the passive component γ_p. It is clear that γ_p is positive (which bounds γ from above by mgh/J), but a lower bound on γ is less clear. Indeed Winter et al. (1998) proposed that γ_p was greater than mgh/J, which would make γ negative and the system marginally stable without active control. While more recent authors have argued against this proposition (Loram and Lakie 2002; Morasso and Schieppati 1999; Morasso and Sanguineti 2002; Peterka 2002), estimates of γ_p have ranged from 91% of mgh/J (Loram and Lakie 2002) to 10% of mgh/J (Peterka 2002). As explained below we constrained our parameter space so that γ lay in the interval $(0, 9)s^{-2}$, and we considered a value of γ less than $0.5\,s^{-2}$ to be unreasonably small.

The other parameters that appear in this plant model, σ, σ_{21}, and σ_{22}, represent noise levels: respectively, the process, visual measurement, and nonvisual measurement noise levels.

Note that in (1), $u(t)$ can only be interpreted as the applied ankle torque if the (2, 1) element of G, $G_{21} = 1/J$. However, as will become apparent below, $u(t)$ can be rescaled by changing the gains on the proportional derivative controller, and these gains are also parameters of the model. To avoid introducing J as a superfluous parameter, we chose $G_{21} = 1$ in (3) and reinterpreted $u(t)$ as the *scaled* applied ankle torque.

We coupled this plant model to four different neural controllers, each utilizing a different adaptive scheme. Each postural model depended upon the four parameters defining the plant (γ, σ, σ_{21}, and σ_{22}) as well as two gains associated with the proportional-derivative control. To compare models, we varied these six parameters to achieve the best fit to experimental data. However, goodness of fit at a particular set of parameter values was determined for all experimental conditions simultaneously. Thus we strove to find a model, together with just one set of values for its parameters, that tested well under many different experimental conditions, without the need for the modeler to adjust parameters between conditions. In testing a successful model only the visual scene position $x_v(t)$ would be under experimental control. In other words, $x_v(t)$ would be the only input to the system. As we have defined the problem, the only output to be compared with experimental data would be the position of the center of mass, $x_1(t)$. The results in this paper were derived from sinusoidal inputs at a range of frequencies and amplitudes. However, in this scenario, stochastic inputs and inputs depending upon the output (i.e., sway referencing) would also be possible.

2.2 Estimating models

The three estimating schemes included progressively more sophisticated internal models of the visual environment. Each of these internal models depended upon parameters (referred to below as *adaptive parameters*) that were allowed to take on a large but finite number of possible values. Each possible set of adaptive parameter values gave a different internal model. The scheme performed state estimation with all of these models separately and simultaneously. In addition, it computed the likelihood that each set of adaptive parameter values correctly described the visual stimulus. The response was based on an average of the state estimates, weighted by their corresponding likelihoods. Thus the modeled nervous system adapted to changes in environmental conditions by recalculating the likelihood-based weights associated with its various internal models.

Each estimating model implicitly assumed that the motion of the environment could be described by an nth-order linear stochastic process of the form

$$\dot{x}_{env}(t) = A\, x_{env}(t) + \xi_{env}(t) ,$$

where ξ_{env} is a Gaussian white noise process with spectral density matrix D. Each scheme assumed that the vector x_{env} had n components, which it interpreted as follows: the first represented the velocity of the visual scene, $\dot{x}_v$, and each subsequent component represented its next derivative. Each scheme expressed A and D in terms of adaptive parameters (d_1, d_2, a_1, and/or a_2). Adaptive parameters are those that the nervous system assumes are unknown or subject to change.

The three internal models of the visual environment differed in how they represented A and D. We called the vector of adaptive parameters θ. For the first estimating model (the random walk model) we assumed

$$A = [0], \quad D = [d_1]^2 .$$

Thus $n = 1$ and $\theta = d_1$. For the second estimating model (the first-order model) we assumed

$$A = [-a_1], \quad D = [d_1]^2 .$$

Thus $n = 1$ and $\theta = (d_1, a_1)^T$. For the last estimating model (the second-order model) we assumed

$$A = \begin{bmatrix} 0 & 1 \\ -a_1 & -a_2 \end{bmatrix}, \quad D = \begin{bmatrix} d_1 \\ d_2 \end{bmatrix} \begin{bmatrix} d_1 & d_2 \end{bmatrix} .$$

Thus $n = 2$ and $\theta = (d_1, d_2, a_1, a_2)^T$. The form we chose for the internal model of the second-order scheme expresses in a canonical way a general linear second-order stochastic differential equation with one observed variable.

To determine the adaptive parameters θ we used a scheme based on likelihood rather than the Myers and Tapley scheme used by Van der Kooij et al. (2001). We made this change because the Myers and Tapley scheme can only be applied to discrete systems – it does not have

a useful continuous time limit. Implementing this likelihood scheme in the dynamic model involved discretizing the adaptive parameter space. Thus we allowed θ to only take on m discrete values:

$\theta \in \{\theta_1, \theta_2, \ldots, \theta_m\}$.

For each possible value θ_k of the estimated parameters the modeled nervous system conducted state estimation under the assumption that the true value of θ was θ_k. In our implementation this process involved m separate Kalman filters. We consider the limit as m becomes large and the discretization becomes dense as the biologically relevant one. While in this limit the dynamic model is intractable to simulate on a serial computer, these computations may be feasible to the human brain given its massively parallel architecture.

Each of the m Kalman filters had a different internal model of the environment as well as the same internal model of the body and sensors. In our implementation, this internal model of the body and sensors was a correct model in the sense that it used the known matrices F, G, Q, H, E, and R from (1) and (2). On the other hand, the internal models of the environment (each involving the matrices $A(\theta_k)$ and $D(\theta_k)$ for some θ_k) attempted to model the unknown stimulus $x_v(t)$ and thus could not be expected to be correct.

To write down the equations for the Kalman filter corresponding to the estimate θ_k, we combine the body and sensor internal model with the environment internal model. Thus the state variables that the nervous system estimates include both the body's state, x, and the environment's state, x_{env}. We write $x^* = (x^T, x_{env}^T)^T$. The nervous system's internal model of the body, sensors, and environment was given by:

$$\dot{x}^*(t) = F^*(\theta_k)x^*(t) + G^*u(t) + \xi_p^*(t)\,,$$
$$z(t) = H^*x^*(t) + \xi_m^*(t)\,,$$

where ξ_p^* and ξ_m^* are independent Gaussian white noise processes with spectral densities $Q^*(\theta_k)$ and R, respectively. These matrices were given by

$$F^*(\theta_k) = \begin{bmatrix} F & 0_{2\times n} \\ 0_{n\times 2} & A(\theta_k) \end{bmatrix}, \quad G^* = \begin{bmatrix} G \\ 0_{n\times 1} \end{bmatrix},$$
$$Q^*(\theta_k) = \begin{bmatrix} Q & 0_{2\times n} \\ 0_{n\times 2} & D(\theta_k) \end{bmatrix}, \quad H^* = \begin{bmatrix} H & E & 0_{2\times(n-1)} \end{bmatrix}, \qquad (4)$$

where the notation $0_{i_1\times i_2}$ indicates a $i_1 \times i_2$ zero matrix. Using this internal model, the Kalman filter associated with the kth possible value of the adaptive estimate, θ_k, maintained an estimate $\hat{x}^*(t, \theta_k)$ of x^*(t). The equations describing the corresponding computations were given by

$$\dot{\hat{x}}^*(t, \theta_k) = F^*(\theta_k)\hat{x}^*(t, \theta_k) + G^*u(t) + K(\theta_k)\left[z(t) - H^*\hat{x}^*(t, \theta_k)\right], \qquad (5)$$

where

$$K(\theta_k) = P(\theta_k)H^{*T}R^{-1} \qquad (6)$$

and $P(\theta_k)$ is the unique symmetric positive definite solution to the Ricatti equation

$$F^*(\theta_k)P(\theta_k) + P(\theta_k)F^*(\theta_k)^T + Q^*(\theta_k) - P(\theta_k)H^{*T}R^{-1}H^*P(\theta_k) = 0\,. \qquad (7)$$

For a derivation of the Kalman filter equations see Bryson and Ho (1975). Bryson and Ho (1975) noticed that the matrices $P(\theta_k)$ and $K(\theta_k)$ did not depend on any dynamic quantities. Our method of solving (7) required that D be positive definite. To enforce this requirement, we restricted our parameter space to the region where $|d_j| > 5 \times 10^{-7}$ deg s$^{-3/2}$ for each j.

In addition to maintaining state estimates for each θ_k, the modeled nervous system also calculated log-likelihoods $q(t, \theta_k)$ of the observed measurements up to time t given that the internal model represented by θ_k correctly modeled the environment. The equation describing this calculation was given by

$$\dot{q}(t, \theta_k) = \hat{x}^*(t, \theta_k)^T H^{*T} R^{-1}\left[z(t) - \frac{1}{2}H^*\hat{x}^*(t, \theta_k)\right] - \epsilon q(t, \theta_k)\,. \qquad (8)$$

For a derivation of this equation [without the term $-\epsilon q(t, \theta_k)$], see Balakrishnan (1973, p. 195). We added the term $-\epsilon q(t, \theta_k)$ as a heuristic to make the system forget past information. Such forgetting was desirable under changing environmental conditions. We called the parameter ϵ the rate constant.

Given the state estimates and the likelihoods, how did the neural controller decide what ankle torque to apply? It computed $u(t)$ in two steps. First it computed an average state estimate $\hat{x}_{av}$ from the state estimates of all the Kalman filters. This average was weighted by the likelihoods

$$\hat{x}_{av}(t) = \frac{\sum_{k=1}^{m} \exp\left(q(t, \theta_k)\right)\hat{x}^*(t, \theta_k)}{\sum_{k=1}^{m} \exp\left(q(t, \theta_k)\right)}\,.$$

Exponentials were used because $q(t, \theta_k)$ represented log-likelihood. The second step based the proportional derivative control on $\hat{x}_{av}$:

$$u(t) = -C^*\hat{x}_{av}, \quad \text{where } C^* = [c_1\ c_2\ 0_{1\times n}]. \qquad (9)$$

Here c_1 and c_2 are, respectively, the proportional and derivative gains, which together with γ, σ, σ_{21}, and σ_{22}, introduced above, comprised the six parameters we varied with the models.

In summary, the stochastic differential equation that defined the estimating models was given by (1), (5), and (8), where (5) and (8) were repeated m times, once for each θ_k. Thus the equation had dimension $2 + m(3 + n)$ where n is the order of the internal model and m is the number of possible discrete values for the adaptive parameter θ.

2.3 Nonestimating model

The three estimating schemes presented above all involved a neural controller that included internal models of the visual environment. These schemes accomplished adaptation to changing environmental conditions by recalculating the likelihood-based weights associated with their various internal models. We now present a postural model that lacks internal models of the environment. We called this scheme *nonestimating* because it did not estimate the dynamics of the environmental motion. Instead it adjusted

sensory weights to minimize the mean square of the control signal without making any specific assumptions about the dynamic properties of the environmental motion.

A second important difference between the estimating and nonestimating schemes was the region of parameter space that they searched for an improved response. The estimating schemes conducted global searches in the sense that they simultaneously compared likelihoods over a wide range of possible adaptive parameter values. In contrast, the nonestimating model made a local search: it considered only the effects of small adaptive parameter changes.

We constructed the nonestimating model by first specifying a linear model with one parameter, then allowing the parameter to slowly vary. Using the same notation as before

$$\dot{x}(t) = Fx(t) + Gu(t) + \xi_{\mathrm{p}}(t)\,, \tag{10}$$

$$\dot{\hat{x}}(t) = (F - GC)\hat{x}(t) + K(\theta)[z(t) - H\hat{x}(t)]\,, \tag{11}$$

where θ is a scalar parameter. As before, $\xi_{\mathrm{p}}(t)$ and $\xi_{\mathrm{m}}(t)$ are Gaussian white noise processes with spectral densities Q and R, respectively, and the measurement vector, $z(t)$, is given by

$$z(t) = Hx(t) + E\dot{x}_{\mathrm{v}}(t) + \xi_{\mathrm{m}}(t), \tag{12}$$

where x_{v} is the position of the visual scene. The matrices F, G, Q, H, E, and R are defined in (3). And finally the control signal u and the matrix C are given by

$$u(t) = -C\hat{x}(t), \quad \text{where } C = [c_1\ c_2]. \tag{13}$$

Note that (11) has the form of a Kalman filter. Many different choices of the matrix-valued function $K(\theta)$ were possible as long as the parameterization allowed the relative weighting of the two measurements, z_1 and z_2, to vary over a wide range. Our choice of $K(\theta)$ was the steady-state Kalman gain matrix for a model of the plant that was correct except in two respects. First, the plant model was incorrect in the sense that it assumed the environment to be stationary: from this assumption came the name *nonestimating*. The second incorrect assumption was meant to mitigate the first incorrect assumption. In particular, the noise level associated with the visual modality was assumed to be θ rather than its veridical value σ_{21}. The idea behind this incorrect assumption was to allow the nervous system to choose an increasingly large value of θ as the amplitude of visual stimulation increased. A larger value of θ implied a plant model that assumed vision to be more noisy – hence one that downweighted vision more. The plant model was correct in every other respect. In particular, the matrices F, G, Q, H, and all other entries of R (besides the upper-left element, σ_{21}) were known and used.

To make these statements precise, $K(\theta)$ could be found by solving the following Ricatti equation for P:

$$0 = FP + PF^{\mathrm{T}} + Q - PH^{\mathrm{T}}\tilde{R}(\theta)^{-1}HP\,,$$

where

$$\tilde{R}(\theta) = \begin{bmatrix} \theta^2 & 0 \\ 0 & \sigma_{22}^2 \end{bmatrix},$$

then using $P(\theta)$ to find $K(\theta)$ as follows:

$$K(\theta) = P(\theta)H^{\mathrm{T}}\tilde{R}(\theta)^{-1}\,.$$

In practice it was simpler to use the following closed-form expression for $K(\theta)$:

$$K(\theta) = \left[\frac{2}{\sqrt{4\gamma + \sigma^2(1/\theta^2 + 1/\sigma_{22}^2)}}\right]\begin{bmatrix} \dfrac{\sigma_{22}^2}{\theta^2 + \sigma_{22}^2} & \dfrac{\theta^2}{\theta^2 + \sigma_{22}^2} \end{bmatrix}. \tag{14}$$

This expression makes apparent the reweighting of the two measurements. When the rightmost matrix factor multiplies z, the result is an average of the two measurements z_1 and z_2, where each is weighted by the square of the assumed noise level corresponding to the other. The leftmost factor then scales this average separately for the position and velocity components.

One interpretation of our choice of $K(\theta)$ is that the nervous system assumes the visual scene to be stationary with all sensory conflicts arising from measurement noise and estimates an appropriate noise level θ to describe this noise. This interpretation is misleading because the scheme does not change θ to improve its estimation of the noise level, whose true value is σ_{21}. Instead it changes θ to minimize the mean square of the control signal. If the visual scene is indeed stationary, θ does not in general converge to σ_{21}.

To derive the equation describing the dynamics of θ, we considered how x and $\hat{x}$ in the linear model [given by (10)–(13)] would change if the parameter θ were changed by a small amount. Differentiating (10) and (11) with respect to θ and using (12) and (13), we obtained

$$\dot{x}_\theta(t) = Fx_\theta(t) - GC\hat{x}_\theta(t), \tag{15}$$

$$\dot{\hat{x}}_\theta(t) = (F - GC)\hat{x}_\theta(t) + K_\theta[z(t) - H\hat{x}(t)] + KH[x_\theta(t) - \hat{x}_\theta(t)], \tag{16}$$

where the subscript θ denotes the partial derivative with respect to θ. The matrix K_θ can be computed by direct calculation from (14). Note that the calculation of $x_\theta(t)$ and $\hat{x}_\theta(t)$ from (15) and (16) involved only the measurement $z(t)$ and other quantities accessible to the nervous system. Thus even though $x(t)$ was not knowable, the neural controller could integrate the equation for its partial derivative with respect to θ, $x_\theta(t)$.

In terms of $\hat{x}_\theta$ we could compute

$$\frac{\partial}{\partial\theta}\left[u(t)^2\right] = 2u(t)u_\theta(t) = 2C\hat{x}(t)C\hat{x}_\theta(t)\,. \tag{17}$$

In the previous calculations we assumed that θ was fixed. We now let θ slowly vary with the goal of minimizing $\langle u^2\rangle$, where $\langle\cdot\rangle$ denotes time average. Equation (17) suggests the gradient descent rule:

$$\dot{\theta}(t) = -\epsilon C\hat{x}(t)C\hat{x}_\theta(t)\,, \tag{18}$$

where ϵ is a small positive parameter that we called the adaptation gain. Simulations showed that (18) did, in fact, tend to minimize $\langle u^2\rangle$ (Fig. 2). (A parameter ϵ also appeared in the estimating models and was called the rate

constant. Its interpretation is different for the nonestimating scheme. Whereas for the estimating schemes the rate constant ϵ determined the rate of forgetting of past sensory information, for the nonestimating scheme the adaptation gain ϵ was related to the rate at which the adaptive parameter changed.)

In summary, the stochastic differential equation defining the nonestimating model was given by (10), (11), (15), (16), and (18). Thus the dimension of the equation was 9. The parameters we varied with the nonestimating model were the same as those we varied with the estimating schemes, that is, we varied γ, σ, σ_{21}, σ_{22}, c_1, and c_2, which defined the matrices F, Q, R, and C in (10)–(13).

2.4 Computing transfer functions

The response to sinusoidal stimuli could be characterized by the transfer function from x_v to x_1. We computed this transfer function, $T_{x_v x_1}$, using two methods. Our first method was the standard one involving the discrete Fourier transform of a simulated model trajectory. Unfortunately, the time required to simulate the model made an extensive search of the parameter space, using this method to evaluate transfer functions, computationally prohibitive. Our second method involved applying an approximate formula for the transfer function, which we derived for each adaptive scheme. This approximate formula was defined in terms of the parameters of the model and the amplitude and frequency of the sinusoidal stimulus. The approximation became accurate in a limit that we state below.

To compute the transfer function, $T_{x_v x_1}$, from simulated trajectories, we calculated the discrete Fourier transforms $X_v(f)$ and $X_1(f)$ of, respectively, $x_v(t)$ and $x_1(t)$, then evaluated both at the stimulus frequency f_s and divided:

$$T_{x_v x_1}(f_s) = \frac{X_1(f_s)}{X_v(f_s)}\,.$$

For each value of the parameters tested and for each experimental condition, we repeated this calculation on ten trajectories each generated with a different seed for the pseudorandom noise. To generate these trajectories, we simulated the model for 2000 s using Euler's method with a time step of 0.005 s, then discarded the first 1000 s to lessen the effect of transients. We computed mean gain and phase as, respectively, the absolute value and angle in degrees of the mean of the ten transfer functions in the complex plane. We obtained a 95% confidence ellipse in the complex plane for the transfer function mean by assuming a bivariate normal distribution. We obtained confidence intervals for gain and phase from bounds on the magnitude and angle of the transfer function derived from the 95% confidence ellipse.

This method for estimating the transfer function required long computer calculations. For the random walk model, with only one adaptive parameter, calculating a confidence ellipse as described above for one set of parameter values and one experimental condition took approximately 1.5 h. The time required for the other estimating models increased exponentially with the number of adaptive parameters. Even for the nonestimating model this method required more than 9 min of computer time. To allow a thorough exploration of the parameter space, we derived a formula that approximated the transfer function of each model. The approximation became accurate in the limit as ϵ (the rate constant or adaptation gain) tended to 0, the transient had decayed, and for the estimating models the discretization was dense in the parameter space. Applying the transfer function formula required less than 1 s of computer time for each model, and simulations verified its accuracy (Figs. 1 and 2).

Three properties common to all four models made these transfer function formulas possible. First, the adaptive parameters tended to approach a steady-state value under the condition that ϵ (the rate constant or adaptation gain) was small. Note that, for the estimating models, what we mean when we say that θ converges to a steady state or asymptotic value is that, as time progresses, only one θ_k in the discretization has a nonnegligible likelihood weight for the computation of $\hat{x}_{av}$ (Appendix B). Second, the asymptotic value of the vector of adaptive parameters coincided with the minimum or maximum of a function that could be explicitly written down and computed in terms of the parameters of the model and the frequency and amplitude of the stimulus. Finally, once the vector of adaptive parameters reached its asymptotic value, the postural system was linear. Thus from the asymptotic value we could compute the approximate transfer function of the model with a straightforward computation.

Unlike computing the transfer function from simulations, there was no pseudorandom element involved in applying the formula. Thus the results of the transfer function formula did not involve confidence ellipses. However, in applying the formula, an optimization step was involved. In consequence, multiple local optima could make the formula fail, depending on the initial condition of the optimization. We mitigated this problem by using multiple initial conditions.

For the estimating models, the asymptotic value of the vector of adaptive parameters approximately coincided with the value $\hat{\theta}$ that maximized the following functional:

$$-\tfrac{1}{2}\mathrm{Tr}\left(R^{-1}H^*\left\langle d_e(\cdot,\hat{\theta})d_e(\cdot,\hat{\theta})^{\mathrm{T}}\right\rangle H^{*\mathrm{T}}\right), \quad \text{where } d_e(t,\hat{\theta}) = x^*(t) - \hat{x}^*(t,\hat{\theta})\,. \tag{19}$$

The angled brackets denote time average. This function could be computed directly from the parameters of the model and the amplitude and frequency of the sinusoidal stimulus. See Appendix B for the derivation of this formula and Appendix C for an explanation of how to calculate it. For later reference we call the functional of (19) the *likelihood correspondence function*. As explained in Appendix B, the estimating models increasingly select the internal model whose parameters maximize this functional. In this sense these systems estimate θ.

Similarly for the nonestimating model, the asymptotic value of the adaptive parameter approximately coincided with the value $\theta^\dagger$ that minimized the functional

$$\left\langle u^2(\cdot,\theta^\dagger)\right\rangle\,.$$

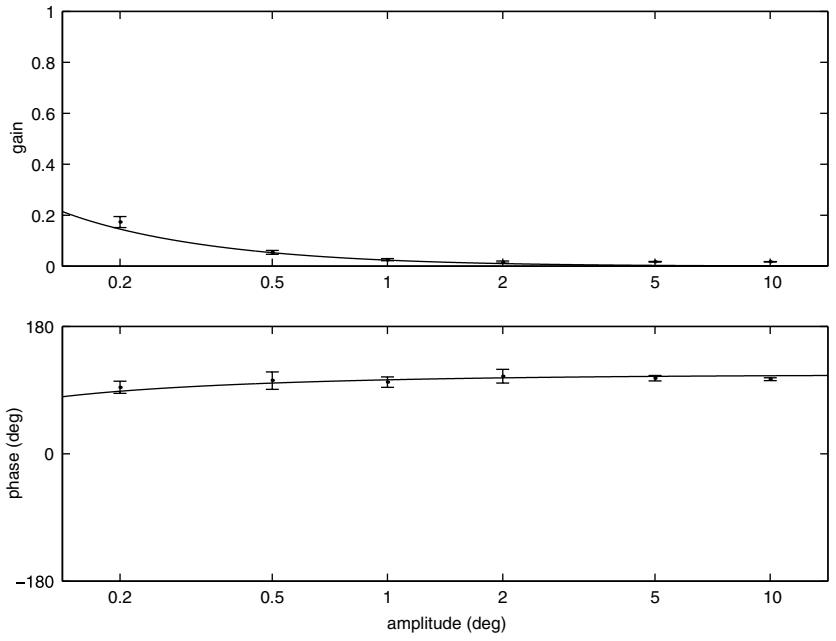

Fig. 1. Comparison of gain and phase calculated by the transfer function formula for the random walk model to estimates of these quantities derived from simulations. The response seen is to 0.2 Hz sinusoidal stimulation as a function of amplitude. The *curves* represent the calculations of the transfer function formula. Gain and phase are, respectively, the absolute value and angle in degrees of the complex valued transfer function. Each *error bar* represents a 95% confidence interval for gain and phase derived from simulations (Sect. 2.4). We used standard parameter values (Sect. 2.6), and for the simulations the rate constant was given by $\epsilon = 0.002\,\mathrm{s}^{-1}$, and the adaptive parameter space was discretized as follows: $d_1 \in \{0.01, 0.05, 0.10, 0.15, \ldots, 2.0\}$

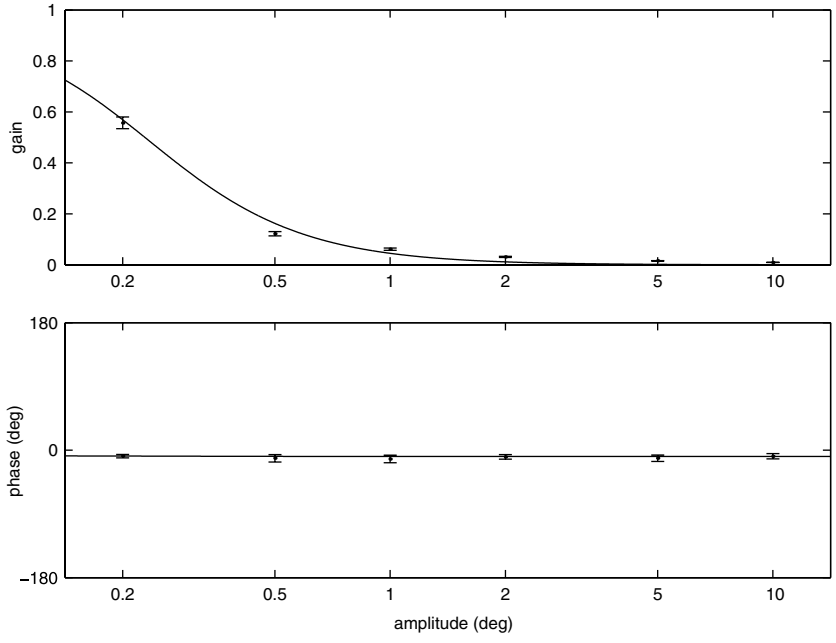

Fig. 2. Comparison of gain and phase calculated by the transfer function formula for the nonestimating model with estimates of these quantities derived from simulations. The response seen is to 0.2-Hz sinusoidal stimulation as a function of amplitude. The *curves* represent the calculations of the transfer function formula. Gain and phase are, respectively, the absolute value and angle in degrees of the complex valued transfer function. Each *error bar* represents a 95% confidence interval for gain and phase derived from simulations (Sect. 2.4). We used standard parameter values (Sect. 2.6), and for the simulations the adaptation gain was given by $\epsilon = 0.002\,\mathrm{s}^{2}$

Again this functional could be computed directly from the parameters of the model and the amplitude and frequency of the sinusoidal stimulus. See Appendix D for an explanation of how to calculate this functional. Note that we have used a different notation for the asymptotic value of the adaptive parameter for the nonestimating model because this asymptotic value could not be considered an estimate, unlike for the estimating models.

In each case, once the asymptotic value of the vector of adaptive parameters was found, the model was linear and its transfer function from x_v to x_1 could be calculated by the method described in Appendix A. Figure 1 compares confidence intervals for the gain and phase of the random walk model derived from simulations with the approximations of the transfer function formula. Figure 2 plots the corresponding comparison for the non-estimating model. The slight discrepancy between the formula and simulations, especially apparent at higher frequencies, is mostly due to the fact that changes to θ happen slowly when θ is large (that is, when the modeled nervous system deems vision unreliable – see Discussion). The dynamics of θ produce very long transients for large-amplitude stimuli. For our parameters and stimulus amplitudes greater than 0.5°, these effects are still significant after 1000 s.

2.5 Search for appropriate response to sinusoidal visual stimulation

We searched the parameter space of each model for the appropriate response to sinusoidal visual stimulation

across different stimulus frequencies and amplitudes. We first attempted a-trial-and-error approach, using different parameters and making many manipulations to the structure of the models. These manipulations included many features not previously described such as sensory dynamics, time delays, position and acceleration coupling, different adaptation times, and state-dependent noise. Only when the unsystematic search for realistic gain and phase behavior failed for each model did we attempt a systematic search. For the systematic search we used the models as described in Sect. 2.1 (i.e., without the manipulations just listed). To allow a thorough exploration of the parameter space, we used the transfer function formulas. Thus we assumed $\epsilon = 0$, the transient was given enough time to decay, and for the estimating models that the discretization was dense in the adaptive parameter space. After restricting each model's (nonadaptive) parameter space, we defined a cost function to quantify goodness-of-fit to experimental data. The domain of this cost function was the tested model's restricted parameter space. The cost was defined in terms of the model's transfer function at different frequencies and different amplitudes of stimulation. We emphasize that all of the transfer functions in a single evaluation of the cost function occur at the same value of the restricted model parameters. Finally, we conducted a systematic search for parameters that optimized our cost function with an algorithm involving simulated annealing.

To restrict each model's parameter space, we assumed that γ was between 0 and $9\,\mathrm{s}^{-2}$. We assumed $c_1 = \gamma + \omega_0^2$ and $c_2 = \alpha$, where ω_0 and α were two of the elements of stochastic structure reported in Jeka et al. (2004) for the fixed surface condition (Table 1). This identification is explained in Kiemel et al. (2002). This restriction left four parameters: $\gamma, \sigma, \sigma_{21}$, and σ_{22}. We then reduced the dimension of the parameter space one further by rescaling the noise terms so that sway standard deviation σ_{COM} equaled the mean sway standard deviation reported in Jeka et al. (2004) for normal subjects with a fixed surface: 0.334°. This rescaling involved initially assuming $\sigma = 1^\circ\,\mathrm{s}^{-3/2}$, calculating σ_{COM}, then dividing each of σ, σ_{21}, and σ_{22} by $\sigma_{\mathrm{COM}}/0.334^\circ$.

For the systematic searches we measured the gain and phase at each of two amplitudes (0.2° and 5°) and at each of three frequencies (0.1, 0.2, and 0.5 Hz). This range of amplitudes and set of frequencies coincided with those studied experimentally in Peterka and Benolken (1995). We imposed the following criteria on the model: phase should not change between the two amplitudes at each frequency, gain should drop by at least 85% between the two amplitudes at each frequency, and, finally, the gain at low amplitude (0.2°) should be between 0.25 and 1.5.

We made these criteria precise with a cost function that depended upon nine measures: three related to changes in phase, three related to gains at low amplitude, and three related to changes in gain. We defined these measures, $m^p_{0.1}$, $m^p_{0.2}$, $m^p_{0.5}$, $m^g_{0.1}$, $m^g_{0.2}$, $m^g_{0.5}$, $m^d_{0.1}$, $m^d_{0.2}$, $m^d_{0.5}$, by

$$m^p_f = \frac{|\Delta p_f|}{64^\circ},$$

$$m^g_f = \begin{cases} \frac{0.25 - l_f}{0.25} & \text{if } l_f < 0.25 \\ 0 & \text{if } 0.25 \le l_f \le 1.5 \\ \frac{l_f - 1.5}{1.5} & \text{if } l_f > 1.5 \end{cases},$$

$$m^d_f = \begin{cases} \frac{\Delta g_f - 0.85\, l_f}{0.85\, l_f} & \text{if } \Delta g_f \ge -0.85\, l_f \\ 0 & \text{if } \Delta g_f < -0.85\, l_f \le 0 \end{cases},$$

where the subscript f denotes frequency ($f \in \{0.1, 0.2, 0.5\}$), Δg_f and Δp_f are the model's changes in, respectively, gain and phase (in degrees) between the two tested amplitudes (0.2° and 5°) at frequency f, and l_f is the gain at low amplitude (0.2°) at frequency f.

We defined the cost function to be the maximum of the nine measures times 100. A cost of zero meant that the model satisfied all our criteria, at that particular value of parameters. The cost could not be less than zero. A cost greater than 100 meant that the phase changed by more than 64° [twice the largest standard error for phase reported by Peterka and Benolken (1995)], the low-amplitude gain was greater than three, or the gain increased rather than decreased over the tested interval. The cost also approached 100 when the low-amplitude gain approached zero. Thus [0, 100] was a typical range for the cost function.

We defined the cost function as a maximum of the measures to insure that the optimization sought to reduce only on the currently largest measure(s). This property insured that at the optimal parameter values the maximal measures could not be simultaneously reduced at the expense of increasing any of the lesser measures.

We performed the optimization with a Nelder–Mead simplex search algorithm combined with a Metropolis (simulated annealing) procedure whereby uphill steps were accepted with a probability that decreased to zero as the optimization proceeded (Press et al. 1992). We repeated the search 50 times starting from parameter values randomly chosen over a wide range.

Table 1. The experimental stochastic structure of quiet stance (means and standard errors for the fixed surface condition) reported in Jeka et al. (2004). Also shown is the models' stochastic structure in quiet stance at the standard parameter values. (The stochastic structure of all four models in quiet stance at standard parameter values agree to three significant digits.)

Measure	Mean	SE	Model
$\beta(\mathrm{s}^{-1})$	0.081	0.028	0.0518
$\alpha(\mathrm{s}^{-1})$	2.188	0.437	2.19
$\omega_0(\mathrm{s}^{-1})$	2.139	0.187	2.17
$\sigma_{\mathrm{COM}}(^\circ)$	0.334	0.038	0.334
$\kappa_r/\kappa_{\mathrm{tot}}$	0.899	0.031	0.889
$2\lvert\kappa_c\rvert/\kappa_{\mathrm{tot}}$	0.136	0.040	0.145

2.6 Standard parameter values

The need to carefully tune parameters to make a model perform realistically (in this case with constant phase across amplitudes) would raise our suspicions if there were no reason to believe that the parameters of the human

postural control system would be so tuned. To compare how well each model performed without parameter tuning we defined *standard parameter values* for each model to be relevant to the behavior of the human postural control system but without reference to the response to sinusoidal stimuli. Specifically, we defined the standard parameter values to be those that best reproduced the stochastic structure of experimental sway in quiet stance. The stochastic structure is a characterization of the dynamical properties of a stochastic process. In particular, a stationary linear process is completely characterized by the autocovariance function of its trajectories. For a linear system, this autocovariance function can, in turn, be completely described by $2n$ numbers, where n is the order of the system. Since these numbers completely characterize the dynamical properties of the linear stochastic process, they comprise its stochastic structure.

To simplify the analysis of the models' stochastic structure, we used their slow-adaptation linear approximation in the quiet-stance condition after the decay of the transient. For a given value of the nonadaptive parameters, all four models had approximately the same stochastic structure in quiet stance.

We used six measures from Jeka et al. (2004) to compare the stochastic structure of the models to that of the experimental data. These measures are defined in terms of the autocovariance function of the sway trajectories:

$$E[x_1(t)x_1(t+\tau)] = \kappa_1 e^{\lambda_1|\tau|} + \cdots + \kappa_p e^{\lambda_p|\tau|}$$

with terms on the right-hand side arranged so that $|\kappa_1 e^{\lambda_1 h}| \geq \cdots \geq |\kappa_1 e^{\lambda_p h}|$, where $h = 0.1\,\mathrm{s}$ is the time step used to analyze the experimental data. The parameters $\lambda_1, \ldots, \lambda_p$ are the eigenvalues of the system. We denoted the first real-valued eigenvalue by λ_r and its corresponding coefficient by κ_r, and we denoted the first pair of complex-conjugate eigenvalues by λ_c and $\bar{\lambda}_c$ and their corresponding coefficients by κ_c and $\bar{\kappa}_c$. Then, $\kappa_r e^{\lambda_r|\tau|}$ is a first-order decay component of the autocovariance function, and $\kappa_c e^{\lambda_c|\tau|} + \bar{\kappa}_c e^{\bar{\lambda}_c|\tau|}$ is a damped-oscillatory component of the autocovariance function. The other terms of the autocovariance function are typically small for experimental trajectories (Kiemel et al. 2002). For the models, we required that $|\kappa_i e^{\lambda_i h}| \leq 0.01\kappa_{\mathrm{tot}}$ for these remaining terms, where $\kappa_{\mathrm{tot}} = \kappa_1 + \cdots + \kappa_p$ is the sway variance. The six measures of Jeka et al. (2004) are (i) the slow-decay rate $\beta = -\lambda_r$, (ii) the damping $\alpha = -(\lambda_c + \bar{\lambda}_c)$, (iii) the eigenfrequency $\omega_0 = \sqrt{\lambda_c \bar{\lambda}_c}$, (iv) the center-of-mass standard deviation $\sigma_{\mathrm{COM}} = \sqrt{\kappa_{\mathrm{tot}}}$, (v) the slow-decay fraction $\kappa_r/\kappa_{\mathrm{tot}}$, and (vi) the damped-oscillatory fraction $2|\kappa_c|/\kappa_{\mathrm{tot}}$.

The standard parameter values for a model were chosen as those that minimized the cost function

$$C = \sum_{i=1}^{6} \left(\frac{m_i - m_i^*}{\delta_i} \right)^2,$$

subject to the constraint that $|\kappa_i e^{\lambda_i h}| \leq 0.01\kappa_{\mathrm{tot}}$ for $\lambda_i \notin \{\lambda_r, \lambda_c, \bar{\lambda}_c\}$, where $m_1^*, \ldots, m_6^*$ and $\delta_1, \ldots, \delta_6$ are the experimental means and standard errors, respectively, of the measures as given in Table 1 and where $m_1, \ldots, m_6$ are the measures for the model. The measures for the model were computed as described in Kiemel et al. (2002).

All four models had the same parameters: γ, c_1, c_2, σ, σ_{21}, and σ_{22}. However, the measurement of velocity and the measurement of velocity relative to the visual scene were indistinguishable in quiet stance. Thus the stochastic structure did not constrain the relative values of σ_{21} and σ_{22}. To make such a constraint we arbitrarily set $\sigma_{21} = \sigma_{22}$.

The stochastic structure of all the estimating models were approximately equal in quiet stance. This implied that standard parameters for all the estimating models were approximately equal. These parameter values were given by $\gamma = 0.19\,\mathrm{s}^{-2}$, $c_1 = 4.9\,\mathrm{s}^{-2}$, $c_2 = 2.2\,\mathrm{s}^{-1}$, $\sigma = 0.25^\circ\,\mathrm{s}^{-3/2}$, $\sigma_{21} = \sigma_{22} = 0.095^\circ\,\mathrm{s}^{-1/2}$. Each standard parameter for the estimating model differed from its corresponding standard parameter for the nonestimating model by less than 2%. Thus the standard parameters for all models were approximately equal. At these parameter values all measures comprising the stochastic structure of the model were within the standard errors given in Jeka et al. (2004) except the slow-decay rate β. At the standard parameter values β was about 36% too small. This discrepancy, as well as the unrealistically low value of γ, may reflect the fact that no computation noise was present in our models (Kiemel et al. 2002).

3 Results

All of the results presented in this section were based on application of the approximate transfer function formulas rather than on simulations of corresponding stochastic differential equations.

3.1 Random walk model

After considerable effort we concluded that the postural control model with the random walk internal model of the visual environment did not behave like the human postural control system in its response to sinusoidal stimuli. Figure 3 shows the gain and phase as a function of stimulus amplitude for three different stimulus frequencies at the standard parameter values. The desired substantial drops in gain came with undesired substantial rises in phase. Using the approximate transfer function formula we found that over the interval [0.2, 5] in amplitude of stimulation the maximum change in phase at the three frequencies tested was 34° at 0.1 Hz. However, at this frequency and over this interval the gain only dropped by 0.13. Greater changes in both gain and phase occurred at lower amplitudes. Specifically, the phase increased by 135° and the gain dropped by 0.65 at 0.1 Hz over the interval [0.1, 5] in stimulus amplitude. Moreover, the model behaved very strangely at low amplitudes. Both gain and phase were initially constant while the estimated parameter d_1 lay at its minimum allowed value: $5 \times 10^{-7}\,\mathrm{deg\,s}^{-3/2}$. At a critical amplitude, which depended on stimulus frequency, all three measures jumped discontinuously.

Changes in the likelihood correspondence function as the amplitude increased explained the first constant then discontinuous behavior of the model. Recall that our adaptive scheme estimated d_1 to be the global maximum

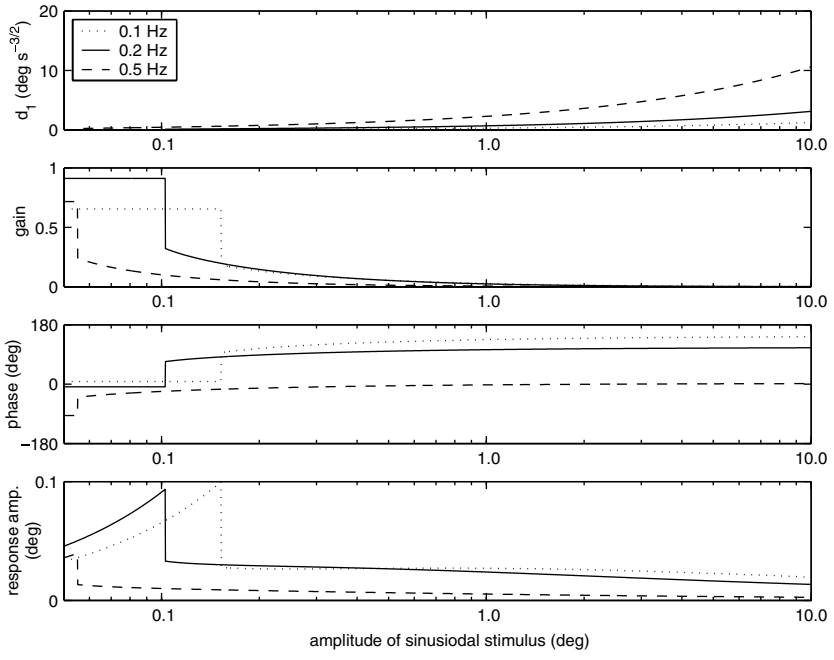

Fig. 3. Estimated d_1, gain, phase, and response amplitude for the random walk model as a function of stimulus amplitude at three different frequencies. Calculations were performed with the transfer function formula. The parameters of the model were those that best reproduced the stochastic structure of postural sway in quiet stance (Sect. 2.6). The response amplitude represents the amplitude of sway at the driving frequency; it is calculated by multiplying the amplitude of the stimulus by the gain. Note that the scale of the response amplitude plot is different here than in Figs. 5 and 6

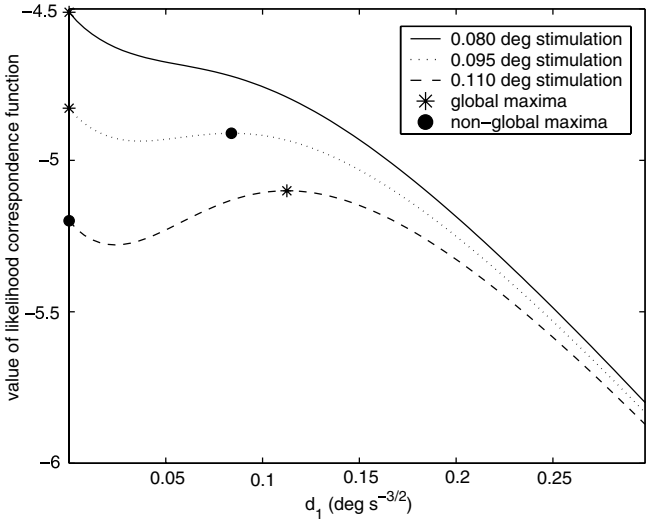

Fig. 4. Likelihood correspondence as a function of d_1 for three different amplitudes of visual stimulation of the random walk model, each at 0.2 Hz

of the likelihood correspondence function. Figure 4 plots this function for three different amplitudes, all at a frequency of 0.2 Hz. If the amplitude was low, its only maximum lay at the endpoint 5×10^{-7} deg $s^{-3/2}$. As the amplitude increased a new local maximum appeared. Initially the new maximum lay below the one at 5×10^{-7}, but at a higher amplitude it overtook the first maximum and became global. At this bifurcation point the estimated parameter, and hence the transfer function, jumped discontinuously.

The postural model with the random walk internal model of the visual environment behaved similarly with different parameter values as well as with many other changes to the structure of the model. We found that it was possible to make the phase nearly constant, as desired, but only at the expense of also making the gain nearly constant, as not desired. The manipulations to the structure of the model that we tried included adding sensory dynamics, coupling the sensory measurements to position and/or acceleration instead of or in addition to velocity, short adaptation time, time delay, and state-dependent noise. This unsystematic search failed to yield realistic behavior. With all these manipulations we consistently observed that with a random walk internal model of the environment significant drops in gain tended to accompany significant rises in phase.

The systematic search produced results that were better, but still not acceptable. At the best-fitting parameters ($\gamma = 0.0022\ s^{-2}$, $\sigma = 0.069^\circ\ s^{-3/2}$, $\sigma_{21} = 0.0032^\circ\ s^{-1/2}$, $\sigma_{22} = 0.23^\circ\ s^{-1/2}$) three measures were maximal (approximately equaling 22.4). Of these three measures, one corresponded to the low-amplitude gain at 0.1 Hz being high (1.83), one corresponded to the low amplitude gain at 0.5 Hz being low (0.19), and one corresponded to the change in phase over the amplitude interval [0.2, 5] at 0.1 Hz being large relative to the criteria quantified by the other measures (14.3°). Thus making the phase relatively constant produced unacceptable gains.

Although this behavior may not seem especially far from the desired behavior, we were not satisfied with these results for two reasons. First, while the model's behavior did lie near the boundary of what we would consider acceptable, it was the best that the model could do for all values of its parameters. In particular, it was the best the model could do without regard to its stochastic structure. For all other values the model would do worse. Thus our optimization procedure found a region in parameter space in which the phase was relatively constant over the interval [0.2, 5] degrees in stimulus amplitude. If the human postural control system implemented the mechanism of the random walk model, there would be no reason to believe

that its parameters would be so tuned. We could see no reason that the human postural control system would evolve to make phase constant over this interval. Mechanisms requiring careful tuning for qualitative agreement with experiment raise our suspicions.

The second reason we were not satisfied with our results is that they were based on a cost function that only looked at the model's behavior for stimulus amplitudes of 0.2° and 5°. The low end of this scale (corresponding to a stimulus of about 5 mm at the eye level of an average person) was high enough to avoid the discontinuous behavior of the model for all frequencies tested. A response to a 2-mm stimulus (about 0.08°) at 0.2 Hz was reported in Oie et al. (2002) without evidence of an abrupt change in the transfer function for larger amplitudes. The random walk model's discontinuity occurs at a larger amplitude for a 0.2-Hz stimulus at standard parameter values. Our failure to make the random walk model behave in a realistic way led us to search for a different mechanism.

3.2 First-order model

A pattern of gain and phase similar to the one seen with the random walk internal model also appeared with the general first-order internal model. Specifically, drops in gain tended to coincide with rises in phase. This tendency is shown in Fig. 5, which plots the gain and phase as a function of amplitude at three frequencies at the standard parameter values. The figure also shows the two adaptive parameters, d_1 and a_1. The constant then discontinuous behavior of the random walk model did not appear. Indeed the likelihood correspondence function, now a function of two rather than one adaptive parameter, had a single local maximum (the global maximum) for all parameters tested. However, we observed that for small amplitudes this global maximum lay on a very broad peak. For small amplitudes all pairs (d_1, a_1) with d_1 small and a_1 had approximately the maximum value of the likelihood correspondence function. In this situation we would expect the weights corresponding to the different Kalman filters to take a long time to converge to their asymptotic values. We did not investigate the consequences of this property on the dynamics of the model.

On the other hand, restricting the likelihood correspondence function to the line determined by the condition $a_1 = 0$ reproduced the likelihood correspondence function for the random walk model. However, neither of the local maxima for the random walk model were local maxima for the general first-order model because likelihood increased as a_1 increased for small d_1.

The systematic search for optimal behavior produced slightly better results for the first-order model than for the random walk model, but these results were still not satisfactory. At the optimal parameter values for the first-order model ($\gamma = 5.5 \times 10^{-5}\,\mathrm{s}^{-2}$, $\sigma = 6.9 \times 10^{-4}\,\mathrm{deg\,s}^{-3/2}$, $\sigma_{21} = 0.0088^\circ\,\mathrm{s}^{-1/2}$, $\sigma_{22} = 0.25^\circ\,\mathrm{s}^{-1/2}$), three measures were approximately maximal (equaling about 19.6). These maximal measures resulted from the low-amplitude gain at 0.1 Hz being high (1.79), the low-amplitude gain at 0.5 Hz being low (0.20), and the change in phase between low and high amplitudes at 0.1 Hz being large relative to the criteria quantified by the other measures (12.5°). Note that the set of measures that were maximal for the first-order model was the same as the set of measures that were maximal for the random walk model, indicating that both mechanisms constrained the postural model in a similar way.

Because the first-order model did not display the discontinuities present in the random walk model, its low-amplitude behavior did not differ substantially from its behavior over the interval [0.2, 5] degrees. Thus tuning the parameters for realistic gain and phase behavior below 0.2° was more successful for the first-order model than for the random walk model. Nevertheless, tuning was necessary to give realistic behavior over this entire range, and even with tuning the tendency for rises in phase to accompany drops in gain could not be removed. As with the random walk model, tuning produced model behavior on the boundary of what we considered acceptable, but we saw no reason why the parameters of the human postural control system would be so tuned. Thus our results cast doubt on the proposition that the human postural control system estimates the state of visual environment with a first-order internal model. This doubt led to a search for a different mechanism that would not lead to a tendency for substantial rises in phase to accompany substantial drops in gain.

3.3 Second-order model

Could a more sophisticated internal model of the visual environment produce more realistic behavior? We found the answer to be a striking no: with a general second- or higher-order linear stochastic process internal model the system responded to any sinusoidal stimulus with a gain that approached zero.

This curious property follows from the fact that a second-order differential equation can correctly model a sinusoidal stimulus. Recall that $x_{\mathrm{env}}(t)$ is the modeled nervous system's representation of the state of the environment. For the second-order model this representation consists of the estimated velocity and estimated acceleration of the visual scene. If

$$A = \begin{bmatrix} 0 & 1 \\ -4\pi^2 f_s^2 & 0 \end{bmatrix},$$

then the solution of the differential equation

$$\dot{x}_{\mathrm{env}}(t) = A x_{\mathrm{env}}(t)$$

with the initial condition

$$x_{\mathrm{env}}(0) = \begin{bmatrix} 2\pi f_s \\ 0 \end{bmatrix}$$

coincides with the true velocity and acceleration of the visual scene $[\dot{x}_v(t), \ddot{x}_v(t)]^{\mathrm{T}}$, provided the stimulus is sinusoidal with amplitude and frequency f_s (more precisely, the true position of the visual scene is given by $x_v(t) = \sin(2\pi f_s t)$). This differential equation is realized as the internal model of the environment with the second-order model if $a_1 = 4\pi^2 f_s^2$, and $a_2 = d_1 = d_2 = 0$.

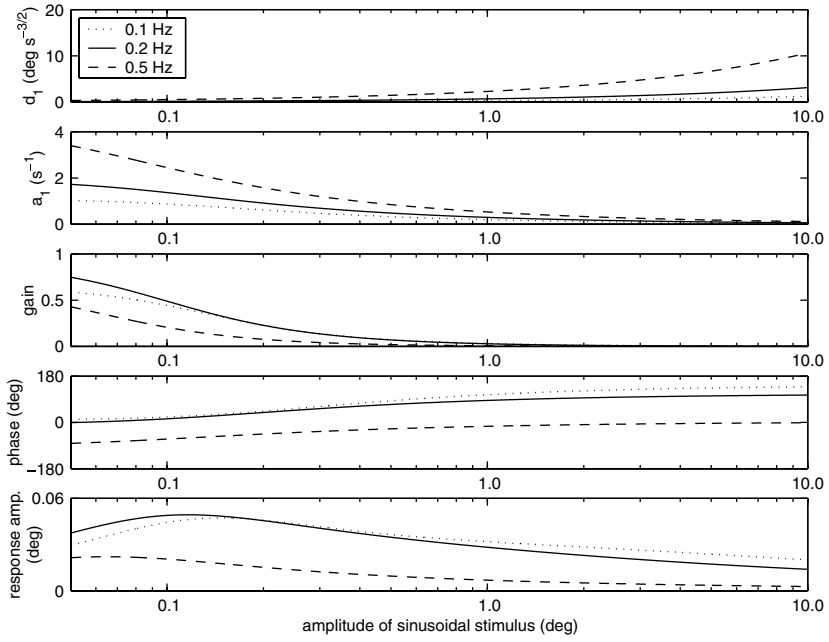

Fig. 5. Estimated parameters d_1, and a_1 and gain, phase, and response amplitude for the first-order model as a function of stimulus amplitude at three different frequencies. Calculations were performed with the transfer function formula. The parameters of the model were those that best reproduced the stochastic structure of postural sway in quiet stance (Sect. 2.6). The response amplitude represents the amplitude of sway at the driving frequency; it is calculated by multiplying the amplitude of the stimulus by the gain. Note that the scale of the response amplitude plot is different here than in Figs. 3 and 6

Not surprisingly, the likelihood correspondence for the second-order model approaches its maximum as

$$[d_1, d_2, a_1, a_2]^{\mathrm{T}} \to [0, 0, 4\pi^2 f_{\mathrm{s}}^2, 0]^{\mathrm{T}} .$$

We called this point in the adaptive parameter space $\theta^{\dagger}$. We found that as θ approached $\theta^{\dagger}$, the transfer function computed by the formula tended to zero.

As θ_k approached $\theta^{\dagger}$, the kth Kalman filter estimated the state of the environment with increasingly high fidelity. (However, this happened only after a transient that decayed slower as θ_k approached $\theta^{\dagger}$.) The modeled nervous system then compensated for the environmental motion by subtracting the estimate from its visual measurement. Moreover, for all amplitudes the Kalman gain remained the same as in quiet stance. Thus the modeled postural control system extracted the same stabilizing information from the visual modality as it did when the visual environment was stationary.

3.4 Nonestimating model

Unlike the models that estimated the state of the environment, the nonestimating model produced realistic gain and phase behavior according to most measures even without tuning the parameters. At standard parameter values, all but one measure was small. Figure 6 shows the gain, phase, and estimated parameter, θ, as a function of amplitude. As evident in the figure, the phase did depend on the frequency of stimulation but did not show a strong dependence on amplitude. Over the interval [0.2, 5] degrees in stimulus amplitude, phase increased by 2.5° at 0.1 Hz, decreased by 0.63° at 0.2 Hz, and decreased by 0.29° at 0.5 Hz. At the same time the gain dropped by more than 85% at each of these frequencies. This behavior, relatively constant phase together with substantial drops in gain, was apparent even without tuning the parameters.

While the match of the nonestimating model to our criteria was good at the typical parameters, it was not perfect. While the gain with a stimulus amplitude of 0.2° was reasonable at low stimulus frequencies (0.7 and 0.6 at 0.1 Hz and 0.2 Hz, respectively) it was only 0.05 at 0.5 Hz. Thus the low-amplitude gains at the lower frequencies were well within our desired range, but the low-amplitude gain at 0.5 Hz was much too low. Nevertheless, tuning the parameters removed this problem. At the parameters that best matched our criteria found by the optimization (approximately $\gamma = 5.7\,\mathrm{s}^{-2}$, $\sigma = 8.2 \times 10^{-6}\,\mathrm{deg\,s}^{-3/2}$, $\sigma_{21} = 0.29^{\circ}\,\mathrm{s}^{-1/2}$, $\sigma_{22} = 0.18^{\circ}\,\mathrm{s}^{-1/2}$), the cost function evaluated to less than the tolerance of the optimization scheme: the phase difference over the tested amplitude interval at all three frequencies was zero to numerical error and the six cost measures related to gain were zero.

Based on the tendency of the estimating models to have phase increases accompanying gain drops, and the lack of such a tendency with the nonestimating model, we favor the nonestimating model over the estimating models.

4 Discussion

We tested four postural models: three containing progressively more sophisticated internal models that estimated the state of the visual environment and one that made no such estimates. We considered the behavior of all four models in response to a sinusoidal stimulus as a function of the stimulus amplitude and frequency. We found that the estimating schemes tended, as a function of increasing stimulus amplitude, to either have substantially increasing phase leads or have no substantial changes in either gain or phase. The nonestimating model, on the other hand, was the only scheme that reproduced the well-known experimental result that across frequencies the gain substantially drops but the phase remains roughly constant as a function of increasing stimulus amplitude.

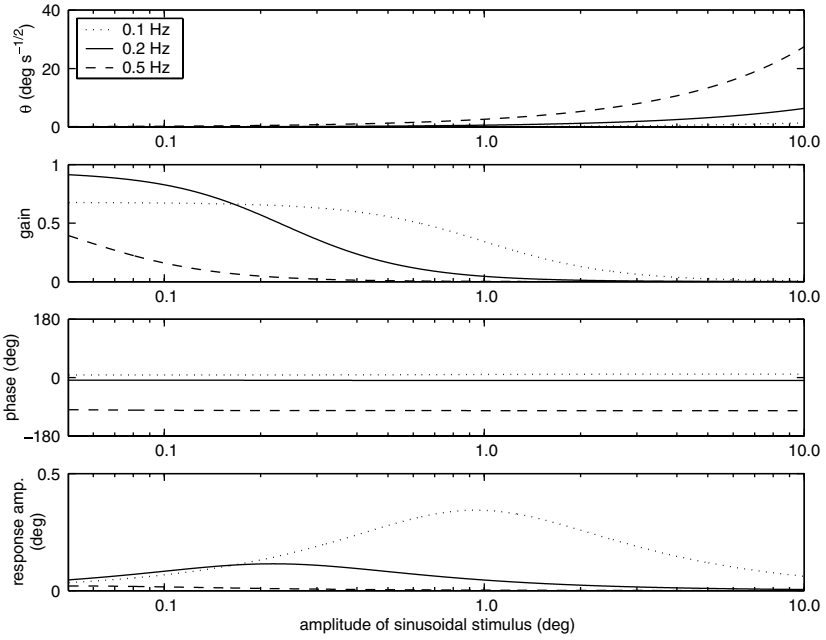

Fig. 6. Asymptotic value of θ, gain, phase, and response amplitude for the nonestimating model as a function of stimulus amplitude at three different frequencies. Calculations were performed with the transfer function formula. The parameters of the model were those that best reproduced the stochastic structure of postural sway in quiet stance (Sect. 2.6). The response amplitude represents the amplitude of sway at the driving frequency; it is calculated by multiplying the amplitude of the stimulus by the gain. Note that the scale of the response amplitude plot is different here than in Figs. 3 and 5

The estimating models, which did not behave appropriately, simulated dynamics of the visual scene while calculating likelihoods corresponding to possible values of the parameters that described the simulations. On the other hand, the nonestimating model, which behaved more appropriately, did not simulate the environment. Instead, the nonestimating model adapted to changing environmental conditions by adjusting sensory weights to minimize the mean square of the control signal. These results suggest that an internal model of the environment is not involved in the nervous system's response to sinusoidal visual stimuli.

Ravaioli et al. (2004) reached a different conclusion concerning the response of the postural control system to translatory visual scene motion. They found that the nervous system extracts stabilizing information from a translating visual scene, and they interpreted this result as evidence that the nervous system uses an internal model's representation of the translation to compensate. The results of Ravaioli et al. (2004) together with ours suggest that the nervous system may handle translatory stimuli differently than sinusoidal stimuli. We leave to future work the task of finding a model of the postural control system that appropriately responds to both sinusoidal stimuli as well as translation-oscillations.

The absence of an internal model in the nonestimating scheme is not the only characteristic that distinguishes it from the other three schemes. The nonestimating scheme uses a local search for changing its adaptive parameter. On the other hand, the three estimating schemes base their estimates on a global search of the adaptive parameter space. In particular, the modeled nervous system with the estimating schemes maintains a separate internal model for each possible set of values of its adaptive parameters. The chosen estimates, upon which the control is based, depend on the calculated likelihoods corresponding to all allowable sets of adaptive parameter values.

The local versus global character of the nonestimating scheme has consequences for the behavior of the system that we consider worthy of a more systematic examination in the future. With the global likelihood scheme, if new values of the adaptive parameters become more likely than the current estimates, the model will immediately adjust. Thus the rate of adjustment to changing environmental conditions will be independent of the current estimates. This property does not hold with the local scheme. In particular, for the local scheme, we have observed that if θ is small (meaning the system considers vision reliable), the system makes changes to θ relatively quickly if conditions change. On the other hand, if this parameter is high (the system considers vision unreliable), it makes changes to θ relatively slowly. The result of this difference is an apparent asymmetry in the system's response to an increase in stimulus amplitude (from low to high) versus a return (from high to low). This statement follows because after an extended period when the amplitude is low, the system deems vision a reliable source of postural information (makes θ small), then changes θ quickly when the amplitude becomes high. On the other hand, after an extended period when the amplitude is high, the system deems vision unreliable (makes θ large), then changes θ slowly when amplitude returns to low. Unfortunately, there is very little experimental data on the dynamics of reweighting to address this prediction.

Is the failure of the estimating schemes a consequence of their global character and not a consequence of their internal models? No. Consider the random walk model. It contains one adaptive parameter d_1 that varies monotonically with stimulus amplitude. In a local scheme, d_1 would have different dynamics, but any reasonable scheme would still estimate a larger d_1 for a larger stimulus amplitude. Gain and phase of the model for large trial times depend on stimulus amplitude only through the asymptotic value d_1 and not through the mechanism through which it is

estimated. We have observed substantial increases in phase accompanying substantial drops in gain as the asymptotic value of d_1 increased with the random walk model. With a local adaptive scheme the random walk model would still have the same property.

The failure of the second-order process model is more interesting. This model is capable of completely compensating for sinusoidal stimuli, extracting the same amount of stabilizing information from the visual modality as when the visual scene is stationary. It shows that the human postural control system does not respond optimally to such stimuli, even though a control scheme that does would evidently not be hard to implement in the nervous system. One can only surmise that sinusoidal stimuli were not common in the environment in which humans evolved. The results of Ravaioli et al. (2004) suggesting the presence of an internal model that compensates for a translating environment may reflect the fact that translating environments were indeed common. It should be pointed out that the second-order process model is only optimal for stimuli that can be described as a linear second-order stochastic process. The model would not behave optimally for a stimulus that consisted of the sum of two sine waves of different frequencies. Such a stimulus would require a fourth-order model for complete compensation.

The results of this paper suggest a different mechanism for adaptation to changing environmental conditions than previously considered: a search for sensory weights to minimize some criterion such as the applied ankle torque. Based on our presented results this model makes one prediction that could prove it false. Note that the response amplitude for the nonestimating model (Fig. 6) rises as a function of increasing amplitude, reaches a maximum, then falls. Likewise, for normal subjects Peterka and Benolken's (1995) data suggest that the response amplitude rises and then saturates as the stimulus amplitude increases. Does the response amplitude also eventually decrease as the stimulus amplitude gets large? For a fixed surface with normal subjects, Peterka and Benolken's (1995) data do not tell a compelling story. If the response amplitude of normal subjects does not decrease, then the nonestimating model, like the second-order model, outperforms the human postural control system in its response to sinusoidal stimuli because it more successfully down-weights the visual modality as it provides more unreliable information. That the human postural control system does not respond in this way would suggest that it was not performing optimally at least with respect to minimizing ankle torque.

Acknowledgements. Funding for this research was provided by National Institutes of Health grant NIH 1RO1NS046065 as part of the NSF/NIH Collaborative Research in Computational Neuroscience Program, John J. Jeka, PI.

References

Arnold L (1974) Stochastic differential equations: theory and applications. Wiley, New York

Balakrishnan AV (1973) Stochastic differential systems: I. Filtering and control. Springer, Berlin Heidelberg New York

Brandt T, Dichgans J, Koenig E (1973) Differential effects of central versus peripheral vision on egocentric and exocentric motion perception. Exp Brain Res 16:476–491

Brandt T, Wist ER, Dichgans J (1975) Foreground and background in dynamic spatial orientation. Percept Psychophys 17:497–503

Bryson AE, Ho YC (1975) Applied optimal control. Hemisphere, Washington, DC

Horak FB, Macpherson JM (1996) Postural orientation and equilibrium In: Shepard J, Rowell L (eds) Handbook of physiology. Oxford University Press, New York pp 255–292

Jeka JJ, Kiemel T, Creath R, Horak F, Peterka R (2004) Controlling human upright stance: velocity information is more accurate than position or acceleration. J Neurophysiol 92:2368–2379

Kiemel T, Oie KS, Jeka JJ (2002) Multisensory fusion and the stochastic structure of postural sway. Biol Cybern 87:262–277

Loram ID, Lakie M (2002) Direct measurement of human ankle stiffness during quiet standing: the intrinsic mechanical stiffness is insufficient for stability. J Physiol 545:1041–1053

Morasso PG, Sanguineti V (2002) Ankle muscle stiffness alone cannot stabilize balance during quiet standing. J Neurophysiol 88:2157–2162

Morasso PG, Schieppati M (1999) Can muscle stiffness alone stabilize upright standing? J Neurophysiol 83:1622–1626

Myers KA, Tapley BD (1976) Adaptive sequential estimation with unknown noise statistics. IEEE Trans Automat Control 22:520–523

Oie KS, Kiemel T, Jeka JJ (2002) Multisensory fusion: simultaneous re-weighting of vision and touch for the control of human posture. Cogn Brain Res 14:164–167

Peterka RJ (2000) Postural control model interpretation of stabilogram diffusion analysis. Biol Cybern 82:335–343

Peterka RJ (2002) Sensorimotor integration in human postural control. J Neurophysiol 88:1097–1118

Peterka RJ, Benolken MS (1995) Role of somatosensory and vestibular cues in attenuating visually induced human postural sway. Exp Brain Res 105:101–110

Press WH, Teukolsky SA, Vetterlin WT, Flannery BP (1992) Numerical recipes in C: the art of scientific computing. Cambridge University Press, Cambridge, UK

Ravaioli E, Oie KS, Kiemel T, Chiari L, Jeka, JJ (2004) Nonlinear postural control in response to visual translation. Exp Brain Res, E-Pub, 9 October 2004

Shumway-Cook A, Woollacott MH (2001) Motor control: theory and practical applications. Lippincott, Philadelphia

Teasdale N, Stelmach GE, Breunig A (1991) Postural sway characteristics of the elderly under normal and altered visual and support surface conditions. J Gerontol 46:B238–B244

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

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

Winter DA, Patla A, Prince F, Ishac M, Gielo-Perczak K (1998) Stiffness control of balance in quiet standing. J Neurophysiol 80:1211–1221

Wolfson L, Whipple R, Amerman P, Kaplan J, Kleinberg A (1985) Gait and balance in the elderly. Clin Geriatr Med 1:649–659

Wolpert DM, Ghahramani Z, Jordan MI (1995) An internal model for sensorimotor integration. Science 269:1880–1882

Woollacott MH, Shumway-Cook A, Nashner L (1986) Aging and postural control: changes in sensory organization and muscular coordination. Int J Aging Hum Dev 23:97–114

Appendix A: Calculation of the transfer function with adaptive parameters fixed

This appendix explains how to compute the transfer function from x_v to x_1, $T_{x_v x_1}$ when the adaptive parameters θ are fixed. (For the estimating models fixed adaptive parameters means $u(t)$ depends on only one state estimate.) With the adaptive parameters fixed all four models are linear.

Consider first the estimating models and fix the adaptive parameters by assuming the kth possible value has significantly greater likelihood than all the others. In this case the kth state estimate will be weighted much more heavily than the other state estimates when $\hat{x}_{av}(t)$ is formed. Thus $\hat{x}_{av}(t) \approx \hat{x}_k(t)$, and we can assume each each estimating model has the form

$$\dot{y}(t) = By(t) + b\dot{x}_v(t) + \xi(t)\,, \tag{20}$$

where $y(t) = (x(t)^T, \hat{x}^*(t,\theta_k)^T)^T$, ξ is a Gaussian white noise process with spectral density S. The matrices B and S and the column vector b are given by

$$B = \begin{bmatrix} F & -GC^* \\ K(\theta_k)H & F^*(\theta_k) - G^*C^* - K(\theta_k)H^* \end{bmatrix},$$

$$S = \begin{bmatrix} Q & 0_{2\times(n+2)} \\ 0_{(n+2)\times 2} & K(\theta_k)RK(\theta_k)^T \end{bmatrix},$$

$$b = \begin{bmatrix} 0_{2\times 1} \\ K(\theta_k)E \end{bmatrix},$$

where F, F^*, G, G^*, C^*, K, H, H^*, E, Q, and R are given by (3), (4), (6), (7), and (9).

Likewise, the nonestimating model, with its adaptive parameter θ fixed, has the same form as (20), but there are slight differences in the definitions of $y(t)$, B, S, and b:

$$y(t) = [x(t)^T, \hat{x}(t)^T]^T$$

$$B = \begin{bmatrix} F & -GC \\ K(\theta)H & F - GC - K(\theta)H \end{bmatrix},$$

$$S = \begin{bmatrix} Q & 0_{2\times 2} \\ 0_{2\times 2} & K(\theta)RK(\theta)^T \end{bmatrix},$$

$$b = \begin{bmatrix} 0_{2\times 1} \\ K(\theta)E \end{bmatrix},$$

where F, G, C, K, H, E, Q and R are given by (3), (13), and (14).

The noise terms do not affect the transfer function from x_v to y. Thus we drop the noise from (20) and compute the Laplace transform of the resulting equation:

$$sY(s) = BY(s) + bsX_v(s)\,,$$

where $Y(s)$ is the Laplace transform of $y(t)$ and $X_v(s)$ is the Laplace transform $x_v(s)$. We now solve for the transfer function from x_v to y at frequency f_s:

$$T_{x_v y} = \frac{Y(i2\pi f_s)}{X_v(i2\pi f_s)} = (i2\pi f_s I_{4+n} - B)^{-1} i2\pi f_s b\,, \tag{21}$$

where I_{4+n} is the $(4+n)\times(4+n)$ identity matrix. (For the nonestimating model $n=0$.)

We are interested in deriving the transfer function from x_v to x_1. To this end, consider the product My where M is a matrix. Then

$$T_{x_v My} = \frac{MY(i2\pi f_s)}{X_v(i2\pi f_s)} = MT_{x_v y}\,. \tag{22}$$

Let $M_1 = [1 \quad 0_{1\times(3+n)}]$ and note $M_1 y = x_1$. Thus $T_{x_v x_1} = M_1 T_{x_v y}$ (that is, $T_{x_v x_1}$ is the first component of $T_{x_v y}$).

Appendix B: Formal derivation of the likelihood correspondence function

This appendix presents a formal derivation of the likelihood correspondence function (19). Throughout we assume that the vector of adaptive parameters θ is fixed in the sense that the $u(t)$ depends on only one state estimate. First we recall the plant model

$$\dot{x}^*(t) = F^*(\theta) + G^*u(t) + \xi_p^*(t)\,, \tag{23}$$

$$z(t) = H^*x^*(t) + \xi_m^*(t)\,, \tag{24}$$

where $\xi_p^*(t)$ and $\xi_m^*(t)$ are independent Gaussian white noise processes with spectral density matrices $Q^*(\theta)$ and R, respectively. Next we recall the equations for the neural computations (5) and (8) and drop the heuristic forgetting term on the equation for $\dot{q}$ (that is, let $\epsilon = 0$):

$$\dot{\hat{x}}^*(t,\theta) = F^*(\theta)\hat{x}^*(t,\theta) + G^*u(t) + K(\theta)\left[z(t) - H^*\hat{x}^*(t,\theta)\right], \tag{25}$$

$$\dot{q}(t,\theta) = \hat{x}^*(t,\theta)^T H^{*T} R^{-1}\left[z(t) - \tfrac{1}{2}H^*\hat{x}^*(t,\theta)\right]. \tag{26}$$

Equations (25) and (26) allow one to compute the log-likelihood for any value of θ. Note that $u(t)$ and $z(t)$ are the same signals for every Kalman filter (that is, for every θ_k).

We now substitute (24) into (26) to obtain

$$\begin{aligned}\dot{q}(t,\theta) &= \hat{x}^*(t,\theta)^T H^{*T} R^{-1} \\ &\quad\times\left[H^*x^*(t) + \xi_m^*(t) - \tfrac{1}{2}H^*\hat{x}^*(t,\theta)\right] \\ &= \hat{x}^*(t,\theta)^T H^{*T} R^{-1} H^*\left[x^*(t) - \tfrac{1}{2}\hat{x}^*(t,\theta)\right] \\ &\quad + \hat{x}^*(t,\theta)^T H^{*T} R^{-1}\xi_m^*(t)\,.\end{aligned} \tag{27}$$

Because ξ_m^* is white noise, the time average of the second term is zero. Therefore, $q(t,\theta)/t$ can be approximated by the time average of the first term. Let $\langle\cdot\rangle$ denote time average. Then for large t

$$\begin{aligned}q(t,\theta)/t &\approx \left\langle \hat{x}^*(\cdot,\theta)^T H^{*T} R^{-1} H^*[x^*(\cdot) - \tfrac{1}{2}\hat{x}^*(\cdot,\theta)]\right\rangle \\ &= -\tfrac{1}{2}\left\langle [x^*(\cdot) - \hat{x}^*(\cdot,\theta)]^T H^{*T} R^{-1} H^*[x^*(\cdot) - \hat{x}^*(\cdot,\theta)]\right\rangle \\ &\quad + \tfrac{1}{2}\left\langle x^*(\cdot)^T H^{*T} R^{-1} H^* x^*(\cdot)\right\rangle,\end{aligned} \tag{28}$$

where the last inequality is obtained using the fact that R^{-1} is symmetric. Note that the second term of (28) does not depend upon θ. Therefore, for some quantity a independent of θ we have that

$$\begin{aligned}q(t,\theta)/t &\approx a - \tfrac{1}{2}\left\langle d_e(\cdot,\theta)^T H^{*T} R^{-1} H^* d_e(\cdot,\theta)\right\rangle \\ &= a - \tfrac{1}{2}\left\langle \mathrm{Tr}\left(R^{-1}H^* d_e(\cdot,\theta) d_e(\cdot,\theta)^T H^{*T}\right)\right\rangle \\ &= a - \tfrac{1}{2}\mathrm{Tr}\left(R^{-1}H^*\left\langle d_e(\cdot,\theta) d_e(\cdot,\theta)^T\right\rangle H^{*T}\right),\end{aligned} \tag{29}$$

where we have used the definition $d_e(t,\theta) = x^*(t) - \hat{x}^*(t,\theta)$ [stated as part of (19)] and the fact that for column vectors v and w $v^T w = \mathrm{Tr}(vw^T)$.

From (29) we recover the likelihood correspondence function (19): for large t the maximum likelihood estimate of θ is approximately equal to the $\hat{\theta}$ that maximizes

$$-\tfrac{1}{2}\mathrm{Tr}\left(R^{-1}H^*\left\langle d_e(\cdot,\hat{\theta})d_e(\cdot,\hat{\theta})^T\right\rangle H^{*T}\right). \tag{30}$$

Observe that a difference in the likelihood correspondence function for two values θ_{k_1} and θ_{k_2} implies a difference in the *rate of growth* of the log-likelihoods $q(t,\theta_{k_1})$ and $q(t,\theta_{k_2})$. Even if the difference in the likelihood correspondence function is small, the difference between the log-likelihoods will grow arbitrarily large as time increases. Thus if there is a value $\theta_{\hat{k}}$ in the discretization for which the likelihood correspondence function is larger than for all other adaptive parameter values in the discretization, its likelihood will grow arbitrarily larger than the likelihoods of all other values. This implies that the weight on the state estimate $\hat{x}(t,\theta_{\hat{k}})$ used to average the state estimates and calculate $u(t)$ will converge to 1 and the weights on all the other state estimates will converge to 0. In this case we say that the control $u(t)$ depends increasingly exclusively on the one state estimate $\hat{x}(t,\theta_{\hat{k}})$ and the system's response becomes increasingly equivalent to the response of the linear model where $m=1$ and the discretization consists of only $\{\theta_{\hat{k}}\}$. In the limit, as the discretization becomes dense in the parameter space, $\theta_{\hat{k}}$ will become arbitrarily close to $\hat{\theta}$. In this sense the system estimates its adaptive parameters.

Appendix C: Calculation of $\langle d_e(\cdot,\hat{\theta})d_e(\cdot,\hat{\theta})^T\rangle$

This appendix presents a method of calculating

$$\langle d_e(\cdot,\hat{\theta})d_e(\cdot,\hat{\theta})^T\rangle, \tag{31}$$

where $d_e(t,\hat{\theta}) = x^*(t) - \hat{x}^*(t,\hat{\theta})$. The method assumes the model's weights have converged so that the one corresponding to $\hat{\theta}$ has approached 1 and the others have approached 0. Using this method to calculate the quantity in (31) coupled with optimizing (30) produces a heuristic method for computing the maximum likelihood estimate of θ. With the weights so converged, (20) gives the model. Now let

$$M_e = \begin{bmatrix} I_2 & -I_2 & 0_{2\times n} \\ 0_{n\times 2} & 0_{n\times 2} & -I_n \end{bmatrix}$$

and observe that

$$d_e(t,\hat{\theta}) = x^*(t) - \hat{x}^*(t,\hat{\theta}) = M_e y(t,\hat{\theta}) + \begin{bmatrix} 0_{2\times 1} \\ x_{\mathrm{env}}(t) \end{bmatrix}, \tag{32}$$

where $y(t,\hat{\theta})$ is the solution of (20). We now use the formula

$$\left\langle d_e d_e^T\right\rangle = \mathrm{Cov}\, d_e + \tfrac{1}{2}^2\Re(\mathrm{T}_{x_v d_e}\bar{\mathrm{T}}^T_{x_v d_e}), \tag{33}$$

where denotes the amplitude of the stimulus, $\mathrm{Cov}\, d_e = E[(d_e - E[d_e])(d_e - E[de])^T]$, $\Re(\cdot)$ denotes the real part of a complex number, and $\bar{\mathrm{T}}^T_{x_v d_e}$ denotes the complex-conjugate transpose of $\mathrm{T}_{x_v d_e}$, the transfer function from x_v to d_e evaluated at the driving frequency $i2\pi f_s$. The drive affects only the mean of d_e and not the covariance. Thus

$$\begin{aligned}\mathrm{Cov}\, d_e &= E[M_e\, y\, (M_e\, y)^T] \\ &= M_e\, \mathrm{Cov}\, y\, M_e^T. \end{aligned} \tag{34}$$

Because the system is linear, the covariance of y is given (Arnold 1974) by the solution of the linear system

$$B\, \mathrm{Cov}\, y + \mathrm{Cov}\, y B^T + S = 0. \tag{35}$$

What remains is to calculate the transfer function from x_v to d_e. Applying (22) to the definition of d_e (32) where $x_{\mathrm{env}} = (\dot{x}_v, \ddot{x}_v, \ldots x_v^{(n)})^T$ we have

$$T_{x_v d_e} = M_e T_{x_v y} + \sum_{j=1}^{n} (i2\pi f_s)^j e^{n+2}_{j+2}, \tag{36}$$

where e^{n+2}_{j+2} denotes the $n+2$ dimensional unit vector with component $j+2$ equal to 1 (and thus all other components equal to 0). The quantity defined in (31) can now be calculated by substituting (34), (35), (36), and (from Appendix A) (21) into (33).

Appendix D: Calculation of $\langle u(\cdot)^2\rangle$

To apply the transfer function formula for the nonestimating model we must instead calculate

$$\left\langle u(\cdot,\theta)^2\right\rangle. \tag{37}$$

Note

$$u(t)^2 = (C\hat{x}(t,\theta))^2 = d_n(t,\theta)d_n(t,\theta)^T,$$

where

$$d_n(t,\theta) = M_n y(t,\theta) = [0_{1\times 2}\ \ C]\, y(t,\theta).$$

To calculate the quantity in (37) we can apply (33) with d_n replacing d_e. For calculating $\mathrm{Cov}\, d_n$ (34) applies (with d_n replacing d_e and M_n replacing M_e). To calculate $T_{x_v d_n}$ we can use (22). Because $d_n(t,\theta)$ is scalar valued, its product with its conjugate transpose is real and is the square of its absolute value. Thus we have

$$\left\langle u^2(\cdot,\theta)\right\rangle = M_n\, \mathrm{Cov}\, y\, M_n^T + \tfrac{1}{2}^2\left|M_n T_{x_v y}\right|^2, \tag{38}$$

where $T_{x_v y}$ is given (in Appendix A) by (21).