

The International Journal of Robotics Research

<http://ijr.sagepub.com>

Perception of Delayed Stiffness

Assaf Pressman, Leah J. Welty, Amir Karniel and Ferdinando A. Mussa-Ivaldi
The International Journal of Robotics Research 2007; 26; 1191
DOI: 10.1177/0278364907082611

The online version of this article can be found at:
<http://ijr.sagepub.com/cgi/content/abstract/26/11-12/1191>

Published by:

 SAGE Publications

<http://www.sagepublications.com>

On behalf of:



Multimedia Archives

Additional services and information for *The International Journal of Robotics Research* can be found at:

Email Alerts: <http://ijr.sagepub.com/cgi/alerts>

Subscriptions: <http://ijr.sagepub.com/subscriptions>

Reprints: <http://www.sagepub.com/journalsReprints.nav>

Permissions: <http://www.sagepub.com/journalsPermissions.nav>

Citations (this article cites 31 articles hosted on the SAGE Journals Online and HighWire Press platforms):
<http://ijr.sagepub.com/cgi/content/refs/26/11-12/1191>

Assaf Pressman*

Sensory Motor Performance Program
Rehabilitation Institute of Chicago
Chicago IL and Department of Biomedical Engineering
Ben-Gurion University of the Negev, Beer-Sheva, Israel
a-pressman@northwestern.edu

Leah J. Welty

Department of Preventive Medicine, Northwestern
University Feinberg School of Medicine, Chicago IL
lwelty@northwestern.edu

Amir Karniel

Department of Biomedical Engineering
Ben-Gurion University of the Negev, Beer-Sheva, Israel
akarniel@bgu.ac.il

Ferdinando A. Mussa-Ivaldi

Sensory Motor Performance Program
Rehabilitation Institute of Chicago and
Department of Physiology, Northwestern University
Feinberg School of Medicine, Chicago IL
sandro@northwestern.edu

Perception of Delayed Stiffness

Abstract

Advanced technology has recently provided truly immersive virtual environments with teleoperated robotic devices. In order to control movements from a distance, the human sensorimotor system has to overcome the effects of delay. Currently, little is known about the mechanisms that underlie haptic estimation in delayed environments. The aim of this research is to explore the effect of a delay on perception of surfaces stiffness. A forced choice paradigm was used in which subjects were asked to identify the stiffer of two virtual spring-like surfaces based on manipulation without visual feedback. Virtual surfaces were obtained by generating an elastic force proportional to the penetration of the handle of a manipulandum inside a virtual boundary. The elastic force was either an instantaneous function of the displacement, delayed at 30 or 60 milliseconds after the displacement or led the displacement (by means of Kalman predictor) by 50 milliseconds. It was assumed that, to estimate stiffness, the brain relates the experienced interaction forces with the amount of penetra-

tion. The results of the experiment indicate a systematic dependence of the estimated stiffness upon the delay between position and force. When the force lagged the penetration, surfaces were perceived as stiffer. Conversely, when the force led the penetration, surfaces were perceived as softer. The perceptual findings were compared with different regression models. This allowed some candidate models to be discarded. To further refine the analysis, a second experiment was carried out in which the delay was introduced only during part of the hand/surface interaction, either while the hand was moving into the spring-like surface or when it was moving out of it. Findings are consistent with stiffness estimates based on dividing the maximum force by the perceived amount of penetration. Findings are not consistent with an estimate of compliance based on the maximum position or local stiffness on the way out nor with linear estimates of stiffness based on the entire force/motion history.

KEY WORDS—Haptics and Haptic Interfaces, Telerobotics, Force and Tactile Sensing, Force Control, Contact Modelling, Virtual Reality and Interfaces

The International Journal of Robotics Research
Vol. 26, No. 11–12, November/December 2007, pp. 1191–1203
DOI: 10.1177/0278364907082611
©SAGE Publications 2007 Los Angeles, London, New Delhi and Singapore

*Corresponding author

1. Introduction

In bilateral teleoperation the human operator can handle remote objects and feel the force feedback in real time (Figure 1a). From a mechanical standpoint, the presence of delays between the operator's movements and the reflected forces introduces apparent active forces when dealing with a passive elastic element. Engineering and theoretical studies of telemanipulation have stressed the importance of stability issues that may emerge from such delays (Adams and Hannaford, 2002; Anderson and Spong, 1989; Hannaford, 1989; Niemeyer and Slotine, 2004).

When we touch an object and try to estimate its rigidity we employ sensory information about the trajectory of our hand and about the interaction force between the hand and the object. This information is processed by the brain's perceptual centers a few hundreds of milliseconds after the event actually occurs (Libet et al., 1979). When we try to touch a remote object through bilateral teleoperation an additional delay is introduced challenging the central nervous system to either change its interpretation of the perceived object or adapt to the additional delay. What is the influence of such delay on the perceived mechanical properties of the manipulated object?

In order to design an effective interface, one should first understand or at least maintain a good hypothesis as to the method by which the brain perceives and handles remote objects. Thus, it was suggested that the brain may employ computations analogous to a Smith predictor (Miall et al., 1993) or wave variables (Massaquoi and Slotine, 1996) for compensating the effects of delays. The ability of the nervous system to adaptively control reaching movements under various external force perturbations was thoroughly explored for state dependent forces (Flash and Gurevich, 1997; Lackner and Dizio, 1994; Shadmehr and Mussa-Ivaldi, 1994) and in some cases for time dependent forces (Conditt and Mussa-Ivaldi, 1999; Karniel and Mussa-Ivaldi, 2003), however the ability to perceive, represent and adapt to delayed force perturbations has not yet been systematically explored. In particular, the influence of delays on the perceived mechanical properties of remote object was largely overlooked. A step toward understanding the influence of delay on the perception of stiffness was taken recently by Hirche (Hirche, Bauer, and Buss, 2005), who evaluated the effect of delay on the transparency of teleoperated systems using wave variables. The work includes a perception study by subjects, which addresses only the presence of change in perception, but not the direction of change.

In this study we considered perception of stiffness of a simple object, a linear spring-like boundary (see Figure 1(b)). We considered various values for the delay, including negative delays, and compared the subjects' behavior to the behavior of a few computational models estimating stiffness by regression, observing extreme force or position and estimating the boundary location. This approach allowed us to address two basic questions: (a) what type of interaction characterizes the explo-

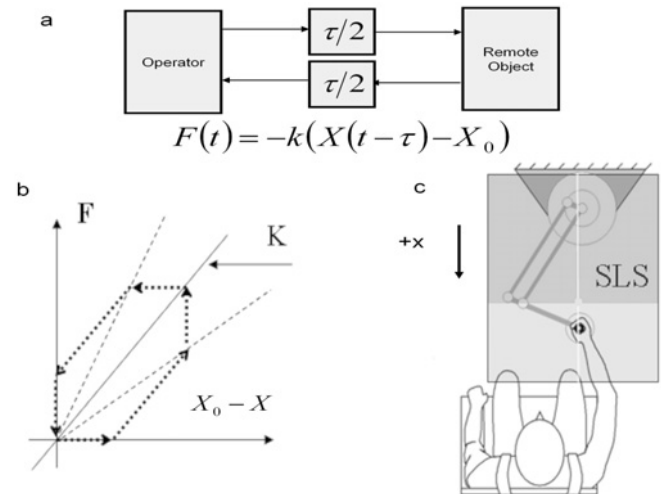


Fig. 1. Teleperception of linear spring-like boundary. **a.** The operator manipulates remote object. If the remote object is a spring, i.e. $F(t) = -k(X(t) - X_0)$, then the operator would experience forces proportional to the position it issued τ seconds before, i.e. $F(t) = -k(X(t - \tau) - X_0)$. **b.** The stiffness of the spring (solid line) and the perception of force during delayed feedback (dotted line). The perception of stiffness could be influenced by the delay. The two dashed lines represent possible stiffness perceptions during the forward and backward part of the movement cycle. With a positive delay (force lagging motion) the perceived stiffness is expected to be smaller than the actual stiffness during the forward push. Conversely, the perceived stiffness is expected to be larger than the actual during the backward motion. The reverse is expected with a negative delay (force leading motion). **c.** A subject holds the robotic manipulandum probing a spring-like surface (SLS) with or without delay. During the experiment, the robotic manipulandum as well as the position of the subject's hand were not visible by the subjects, who saw only the projected spring-like surface and a vertical line indicating the location of their hand along the Y axis. A bright point was projected at a fixed location and the subject was instructed to keep the line near this point.

ration and estimation of an elastic boundary? and (b) how do delays in a teleoperated system affect the perception of an object's rigidity? In the next section we describe the experimental setup and protocol. Then, we report results that address these questions based on comparing the outputs of different estimation models with the judgments of five subjects that probed a spring-like surface with and without delay. We use this comparison for ranking the plausibility of the estimators as computational models of the neural processes that evaluate force and motion information for determining the rigidity of a contact.

2. Methods

2.1. Subjects, Apparatus and Protocol

Five subjects (three male and two female) participated in this experiment after signing the informed consent form approved by Northwestern's Institutional Review Board. Seated subjects held, with their dominant (right) hand, the handle of a two-degree of freedom robotic manipulandum, and looked at a screen, placed horizontally above their hand, which displayed in alternation one of two virtual spring-like surfaces as green or red wide squares (Figure 1c). For further details about the robotic manipulandum, see (Mussa-Ivaldi, Hogan, and Bizzi, 1985; Shadmehr and Mussa-Ivaldi, 1994). The robotic manipulandum exerted forces on the subject's hand and acquired its trajectory. The location of the hand was displayed by a line perpendicular to the boundary of the object. This provided subjects with partial position information, which included the lateral position of the hand without revealing the degree of penetration inside the virtual object. By keeping this line at the same location, subjects contacted the objects at fixed configuration of the arm. The task required that subject probe the object rigidity by executing back and forth motions of the handle within the object's boundary. The experiment was based on a forced choice paradigm: in each trial the subject was presented with two surfaces: one in which the stiffness was varied across trials (surface *K*), and the other in which the delay between the force and displacement was varied across trials (surface *D*). The two surfaces were represented by rectangles of different colors – red and green. The two colors however were assigned randomly, so that each surface type (*K* or *D*) was not uniquely associated with a color. Whenever the hand of the subject moved out of a surface by a fixed amount, the object's switched between *K* and *D* types and the display changed color accordingly. Subjects could switch between the two surfaces as many times as they pleased until they felt ready to answer the question: "Which surface is stiffer (green or red)?" The answer was given by the subject pressing one of two buttons on a custom made hand held device. No feedback was provided after each trial. Three types of data were acquired:

1. The responses of the subjects (green or red, for each trial).
2. The position of the hand along the *X* axis, sampled at a rate of 100 samples per second. Position data were used in real time to generate the virtual surfaces.
3. The interaction force with the surface. This force was calculated in real time according to the hand position, the delay and the elastic properties of the surface (stiffness and boundary).

The force exerted by the virtual surfaces was in the *X* axis direction (see Figure 1c), in proportion to the displacement from the boundary, X_0 , i.e.

$$F_x(X) = \begin{cases} -k(X(t - \tau) - X_0) & \text{if } X < X_0 \\ 0 & \text{if } X \geq X_0 \end{cases} \quad (1)$$

Where F_x is the force in the *X* axis direction, k is the spring's stiffness constant, $X(t)$ is the position along the *X* axis and X_0 is the coordinate of the boundary. During each of the 400 trials presented to the subject, the *K* surface took one of the stiffness values drawn randomly from 150 to 600 N/m (in increments of 50 N/m). The stiffness of the *D* surface was set to 375 N/m, therefore the stiffness of the *K* and *D* was never equal. The *D* surface was assigned delay values, τ , drawn randomly from -50, 0, 30, 60 ms. The delay of the *K* surface was set to 0.

For causal impedance, positive values of τ , would cause the force to lag τ seconds after the position. A situation in which the force is leading the position, i.e. with negative τ , can be regarded as "non-causal". Such a non-causal surface was emulated for relatively short values of τ . For these short values, and the presence of low measurement noise, time derivatives were estimated with relatively small error in an optimal manner using a Kalman filter (Kalman and Bucy, 1960; Grewal and Andrews, 2001). Using these derivatives (up to the 2nd order), the slow varying hand trajectory signal was predictively estimated using linear regression (Chatterjee and Hadi, 1986).

By keeping the hand at rest "inside" the surface, one would eventually eliminate the effects of delay. Therefore subjects were asked to keep the hand in motion while "inside" the probed surface. In each trial the subjects performed numerous surface contacts with the *K* and *D* surfaces, at different times. A typical trajectory is presented in Figure 2 for such contacts.

In a second experiment, a more complex, nonlinear stiffness was used, in which the delay took different values on the inward and outward portions of the probing movement. The force during the inward probing motion could have zero delay while on the outward motion the force was lagging the position ($\tau > 0$), or vice versa, the force lagging the position on the way in and not delayed on the way out (see Figure 4a and 4c for examples). Inward motion is a hand movement with negative velocity and outwards motion has positive velocity (Figure 1).

During this experiment five subjects (who also participated in the first experiment) were presented with 300 trials randomly chosen from the following groups:

- No delay.
- Delay of 50 ms on the way into the surface and zero delay on the way out.
- Zero delay on the way into the surface and 50 ms delay on the way out.




Fig. 2. Typical trajectory during the probing of a virtual surface. The different line styles represent the surface contacts with the K surface (solid) and D surface (dashed). Each contact motion was typically composed of elementary in-out movements in which the hand reached a maximum penetration. A single surface probe is indicated by the gray rectangular region. The gray horizontal line represents the edge of the surface, i.e. X_0 in Equation 1.

The stiffness of the D and K surfaces was 375 N/m and 150 to 600 N/m (in increments of 50 N/m) respectively, the same as in the first experiment.

2.2. Data Analysis

Psychometric curves were derived by estimating the frequency with which subjects indicated that the D surface was stiffer than the K surface, as a function of the actual stiffness difference $K_D - K_K$. Figure 3 show hypothetical psychometric curves describing subject replies as a function of the gap between the stiffness of the surfaces D and K .

For two linear spring-like surfaces without delay (i.e. for delay of the D surface equal to zero) we expect the answers to reach chance level when the two values of the stiffness are equal. A sigmoidally shaped curve is expected, as answers are likely to be correct when the stiffness levels are substantially different.

Specifically, let K_D be the stiffness of the D surface and K_K the stiffness of the K surface, and let Δ be the difference between these two stiffness values. Each point on the psychometric curve in Figure 3 is an estimate of the probability of the subject reporting that the D surface is stiffer than the K surface as a function of the actual stiffness:

$$\hat{P}(\text{Reporting that } \hat{K}_D > \hat{K}_K) = \frac{\sum_{j=1}^{R(\Delta)} N^j}{R(\Delta)} \quad (2)$$




Fig. 3. Possible psychometric curves of the expectation of an answer indicating that D is stiffer than K . The X axis is the difference in stiffness between surface D and K . The black dot-dashed line demonstrate the performance of a “perfect subject” who can accurately estimate whether surface D is stiffer than K . The black solid line shows the “regular subject”, which would make some mistakes in the transition region (marked as a gray rectangular). A shift in this graph to the left(right), as seen by the dashed black line on the left(right), would suggest the subject perceives surface $D(K)$ as stiffer than it really is.

where

$$N^j = \begin{cases} 1, & \text{reply } \hat{K}_D > \hat{K}_K \\ 0, & \text{reply } \hat{K}_D < \hat{K}_K \end{cases} \quad (3)$$

j is an index over the trials with stiffness difference Δ . $R(\Delta)$ is the number of times this stiffness difference was encountered during the whole experiment. The “hat” sign over the stiffness of the K and D surfaces denotes the perceived stiffness by the subject, as opposed to the actual nominal stiffness of the surface. The psychometric curves were calculated separately for each value of τ that was introduced in the D surface. The stiffness at the crossover point (where $\hat{P} \approx 0.5$) is evaluated using a maximum likelihood fit of a sigmoid function to the data points. The Bootstrap method was used in order to estimate the goodness of fit (Wichmann and Hill, 2001b; Wichmann and Hill, 2001a).

2.3. Nominal and Perceived Penetration

We define the “nominal penetration” as the distance traveled from the spring’s nominal zero point (see Equation 1) to the

point of maximum penetration. These distances are equal when no delay is present between position and force. The presence of delay causes a change in the effective location of the boundary of the surface – i.e. the location at which the subjects begin experiencing an elastic force. Therefore, the actual point in which the force effectively changes from zero is different from the nominal zero point.

Indicating by $P_0(k)$ the average penetration with stiffness k and zero delay and with $\delta(k, \tau)$ the average displacement of the boundary with stiffness k and delay τ , the average nominal penetration for the same stiffness and delay values is modeled as:

$$P_{nom}(k, \tau) = P_0(k) + \alpha\delta(k, \tau), \quad (4)$$

This linear relation is supported by experimental observations (Figure 9). The value of α was estimated from all trials in which the subject was probing a surface where the delay (τ) was different from zero. For each trial P_{nom} (the maximum displacement, see Figure 4b) and δ (the amount the surface was shifted due to the delay) were measured and are held in the vectors \vec{P}_{nom} and $\vec{\delta}(k, \tau)$, respectively. For surfaces with zero delay the value of $\delta(k, 0)$ is zero. These trials were used to estimate P_0 by averaging the amount of penetration across all trials. The value of α was estimated in the least square sense by using:

$$\begin{aligned} \vec{P}_{nom}(k, \tau) &= P_0(k) + \alpha\vec{\delta}(k, \tau) \\ \Rightarrow \alpha &= \left(\hat{P}_{nom}(k, \tau) - P_0 \right) \cdot \hat{\delta}(k, \tau)^+, \end{aligned} \quad (5)$$

where $\hat{\delta}(k, \tau)^+$ is the pseudo inverse of $\hat{\delta}(k, \tau)$ (Golub and Loan, 1996).

If one makes the assumption that the perceived penetration corresponding to a given stiffness remains invariant across delays and equal to the zero delay penetration, i.e.

$$P_{per}(k) = P_0(k), \quad \forall \tau, \quad (6)$$

then the perceived penetration is derived from the nominal penetration and the boundary displacement as

$$P_{per}(k) = P_{nom}(k, \tau) - \alpha\delta(k, \tau) \quad (7)$$

where the dependencies on τ on the right side cancels.

2.4. Models

In order to try to explain the behavior of the subjects for various delays, we compared the subjects' responses to the responses generated by different computational models. Each model formulated a comparison of rigidity between the same pairs of surfaces that were evaluated by the subjects. Model estimates were based on force and position data collected during each experiment. Figure 4b shows a force/position trajectory




Fig. 4. **a** Delay on the way out. The slope of the gray slant line is the estimation of the stiffness using the F/P (the estimation of the P_{per} model for this case is the same). Note that the time course of the motion is counter clockwise, due to the positive delay. **b** A single surface probe in the force displacement plan with a constant delay and the parameters extracted: Maximum displacement (large gray circle), peak force (large black circle). The slope of the gray solid line represents an estimation of stiffness carried over the whole penetration data. Local stiffness portion of data from the way in (black circles) and the estimated slope of that portion (gray dashed line). **c** The slope of the gray solid (gray dashed) line is the estimation of the stiffness using the F/P (F/P_{per}) model.

derived from a single probing motion into the virtual surface, in the presence of a delay. The model responses were compared with the subjects' responses. All models have the property that they would produce a correct response when comparing surfaces with different stiffness and no delay. In more detail the models are as follows:

Global stiffness (#1): The model estimates a surface's stiffness by linear regression on the entire set of force and position data. The slope of the gray solid line in Figure 4b is an example.

Maximum Force (#2): The model uses the maximum force exerted by the virtual surfaces on the subject as a proxy for stiffness, assuming that higher forces from either the K or D surfaces would lead to a report of higher stiffness. The black large circle in Figure 4b corresponds to the value of maximum force.

Maximum displacement (#3): The model uses the maximum penetration into the surfaces as proxy for compliance (inverse of stiffness). Thus, deeper penetration into either the K or D surfaces, would lead to a report of lower stiffness. The large gray circle in Figure 4b corresponds to the value of maximum penetration.

Peak-Force/P ratio (#4): The model estimates the stiffness by dividing the maximum force by the amount of penetration, measured from the point where the force is different than zero to the point where maximum force has occurred. The slope of the gray solid line on Figure 4c is the estimated value.

Peak-Force/P_{per} ratio (#5): The model essentially is the same as the Peak F/P ratio model, but the normalizing factor is adjusted according to equation 7, to account for the bias due to the subject’s estimate of the surface boundary position. The slope of the gray dashed line on Figure 4c is the estimated value.

Local stiffness way out (#6): In local stiffness models, the stiffness is estimated based on a portion of the force and position data. In this model, the data from the point of maximum penetration up until exiting the surface is used. Labeling the samples of the force data as *F* and the samples of the position data as *X*, stiffness might be estimated as $\hat{k} = F \cdot X^+$ where the X^+ stands for the pseudo-inverse of *X* (Golub and Loan, 1996) and accomplishes a least square estimation for the value of \hat{k} .

Local stiffness way in (#7): similar to model #6 using data points from entering the surface until maximum penetration. These are the black circles on Figure 4b.

Reference Model (#8): This is an upper limit for the possible model performance assuming normal distribution of the errors. We refer to this as Model #8 and describe its details in subsection 3.6 and Figure 5.

During each trial several probing motions were made by each subject. To derive a single value for each trial, the median across the estimated values in each trial was taken.

A binary comparison between the output of each model and the response of the subject, on a trial by trial basis, was used to estimate the ability of each of the models to predict the behavior of the subjects. The detailed procedure is explained in the next section.

2.5. Model simulation

On the *i*th trial the subject generates the force and position data set $Z^i = \{F(t)_D^i, X(t)_D^i, F(t)_K^i, X(t)_K^i\}$ where the superscript represents the trial number and the subscript the position and force data gathered during the probing of the *K* or *D* surface (“*K* data” and “*D* data”). We use a forced choice paradigm, where the response to the question “Which surface is stiffer?” is expressed by the binary variable $N \in \{0, 1\}$. $N = 1$ corresponds to the subject response of surface *D* being stiffer than surface *K*. Each model was regarded as an operator $M_q(\cdot)$ where the subscript *q* ($q = 1 \dots Q$, *Q* = the number of models) is the model number. We indicate the output of the *q*th model on the *i*th trial as $\tilde{N}_q^i = M_q(Z^i)$ where $\tilde{N}_q^i \in \{0, 1\}$ and Z_i is the data set for the *i*th trial. The \sim symbol over the *N* stands for the prediction made by the *q*th model




Fig. 5. The formalization of the reference model. The top left panel illustrates the probability density function for the *K* surface. The different shades of gray correspond to different stiffness encountered. The bottom left panel depicts the probability density function for the *D* surface, which is shifted by an amount of μ_D . The panel on the right illustrates the outcome frequencies in which the comparison between the two distributions will report the *D* surface being stiffer than the *K* surface. The dashed and dashed-dotted lines show how the two surfaces are compared. The response to the question “is the *D* surface stiffer than the *K* surface?” when the stiffness value of the *K* surface is drawn from the light shaded Gaussian (on the left of the top left panel - dashed gray arrows) is likely to be 1 (yes). On the other hand a sample from the most right distribution will most likely be regarded as 0 (dashed-dotted arrows), since the *K* surface is probably stiffer than the *D* surface.

for the binary value of the outcome. The perfect model will have:

$$\tilde{N}^i = N^i \quad \forall i \tag{8}$$

The best model is the model that generates the best match with the subjects’ responses in all the experimental conditions. Each model operates twice, once on the *K* data and once over the *D* data. Let $E_{D,q}^i = E_q(X(t)_D^i, F(t)_D^i)$ be the estimation over the *D* data (*i*th trial) by model *q* of the associated property (i.e. global stiffness for model #1, local stiffness for model #6 or #7, maximum force for model #2, and so on) and let $E_{K,q}^i = E_q(X(t)_K^i, F(t)_K^i)$ be the estimation of the same property by the same model using the *K* data (*i*th trial). The output of each model can be described by the two outcomes as follows:

$$\tilde{N}_q^i = M_q(Z^i) = \begin{cases} 1, & \rho_q^i < 0 \\ 0, & \rho_q^i \leq 0 \end{cases} \tag{9}$$

where, ρ is defined as:

$$\rho_q^i = E_{D,q}^i - E_{K,q}^i \quad (10)$$

2.6. Subject responses as reference model

A statistical model directly derived from the responses given by the subjects provides an upper limit for the best achievable performance by the other models. In this reference model (regarded as model #8), the estimators for the K and D data were assumed to be normally distributed, with mean μ and standard deviation σ , that is $\hat{K}_D \sim N(\mu_D, \sigma_D)$ and $\hat{K}_K \sim N(\mu_K, \sigma_K)$. This implies that the “perceived” stiffness by the model of each surface is randomly distributed around some mean value. A delay between the force and position might change the value of μ_D to reflect the change in perception, and it might also change the spread of the distribution (σ_D). One should note that the reference model utilizes two different estimators, for the D data and the K data.

Assuming that the perceived shift in stiffness by a subject (derived by using the subject’s data, equation 2 and maximum likelihood fit of a cumulative Gaussian function (Wichmann and Hill, 2001b; Wichmann and Hill, 2001a)) is the mean value for the \hat{K}_D estimator and that the nominal stiffness of the K surface is the mean for \hat{K}_K estimator, a predication of the subjects reply can be made on each trial. The value of the standard deviations (σ_D and σ_K) was set based on the slopes of the psychometric curves for all the subjects and was manually set to 20 N/m. It is presumed that a psychometric curve, generated by the predication in each trial (such as equation 2 would generate) would have 50% in the vicinity of μ_D .

Figure 5 illustrates both the probability density function of the \hat{K}_D and \hat{K}_K estimators outcomes (left panels) and the result of equation 2 for the outcome reported by the two estimators (right). The top left panel shows the reported stiffness for each one of the stiffness values encountered during the experiment for the K surface (for ease of drawing only six values are plotted and not ten). Each Gaussian shade notes a different stiffness level of one of the K surfaces. The panel shows the *density function* of the distribution; during each trial only one of the samples from a single Gaussian is compared to a single sample creating the density function of the D surface’s distribution. The bottom left panel displays the distribution of the reported stiffness for the D surface. It is apparent that the distribution is shifted from the nominal stiffness of the surface (K_D thin dotted black vertical line) by an amount μ_D (thick dashed black vertical line). The right panel shows the result of applying equation 2 to the two distributions (and using equation 9). Since equation 2 reports the estimated probability of the D surface being perceived as stiffer than the K surface, a trial in which the stiffness is drawn from the darkest distribution of the K surface (the most right in the top left panel) is likely to be perceived as stiffer than the D surface (follow

the two gray the two gray dot-dashed arrows on the figure). In such a comparison the perceived stiffness of the D surface is smaller than the perceived stiffness of the K surface; therefore $K_D - K_K$ is smaller than zero. The probability of the subject reporting that the stiffness of the D surface is larger (this is actually the function on the right panel) is close to zero. Therefore this comparison would be mapped to the left portion of the estimated probability function on the right panel. On the other hand, a trial in which the stiffness is drawn from the lightest distribution (far left, top panel) would produce a probability that tends to one.

One has to be clear about this statistical model and its difference from the models described in the previous section. The latter are entirely deterministic and make use of kinematic and force data acquired from subjects motions and interactions forces to determine the relative difference in stiffness of the D and K surfaces. In contrast, the parameters μ and σ used for the statistical model are estimated from the *responses of the subjects*. Therefore the statistical model only serves as a reference estimate of the relative difference in stiffness between the D and K surfaces, based on subjective responses. It does not provide any idea about how these responses were derived from sensory motor information.

3. Results

3.1. Stiffness was overestimated when force lags motion and was underestimated when force leads motion

The psychometric curves from one subject are shown in Figure 6 and clearly demonstrate that delayed stiffness for positive values of τ (the delay of force with respect to position in equation 1 is estimated as being larger than non-delayed stiffness ($\tau = 0$)). This is evidenced by the shift of the crossover point toward values of delayed stiffness that are smaller than the non-delayed stiffness (i.e. $\Delta \equiv K_D - K_K < 0$). Accordingly, subjects tended to equate the value of the non-delayed stiffness to a lower value of delayed stiffness. Consistent with this, for negative delays, the stiffness was under-estimated. This is clearly not an artifact induced by uncontrolled factors in the paradigm and apparatus, as the baseline curve (black line, triangular markers) crossed correctly the chance level at 0 stiffness difference. The same effects were observed in all five subjects. It is interesting that the slope of the fitting sigmoid decreased systematically with increasing values of delay. This trend was observed in all subjects and reflects an increasing difficulty to assess the stiffness of the surface as the delay of the reflected force increases. One could note that the fitting seems to ignore several data points that are not in line with all the other points, e.g. the diamond of the -50 ms delay curve. These samples are referred to as lapses (Wichmann and Hill, 2001a) and are the outcome of stimulus-independent errors made by the subject. Such points might bias the estima-




Fig. 6. The average reply to the question is plotted as function of the difference between the stiffness of the D surface and the K surface $K_D - K_K(j)$, where K_D is the stiffness of the surface with delay, K_K is the stiffness of the surface with variable stiffness and $j = 1 \dots M$, M is the number of different stiffness during the experiment. The points in various shades are the results for the different delays (as they appear in the legend). Each point indicates the number of times the subject assumed the D surface to be stiffer than the K surface, divided by the total number of times that particular stiffness difference was encountered. The lines of various shades are maximum likelihood fits of sigmoid functions to the points. The horizontal lines on each sigmoid represent the confidence level in the position along the X axis at 0.5 detection probability.

tion of the parameters that describe the psychometric curve. In order to reduce the bias caused by these lapses, the fitting procedure does not assume that the upper and lower bound of the psychometric curve are one and zero respectively. Alternatively, the assumption is that these two values are drawn from two narrow uniform distributions: 0–0.05 for the lower bound and 0.95–1 for the upper bound. The outcome is a curve which might not reach the upper and lower bound, or pass through each point, but, on average, would better describe the data.

3.2. Ranking of the models: Three models (#1 Global stiffness, #3 Max displacement and #7 Local Stiffness on the way out) were found to be inconsistent with the subjects' behavior

Agreement between a model and the subject on some i th trial is achieved once both the subject and the model report the same answer to the question “Which surface is stiffer?”, regardless of the actual stiffness difference of the two surfaces, as stated




Fig. 7. The probability of agreement for the various models over all delays for the first experiment. The black dots are the mean value across subjects and the wings represent 95% confidence intervals. The WI and WO on the Local Stiffness estimators stand for way in and way out, correspondingly.

in equation 8. Figure 7 displays the mean of the “probability of agreement” $P(N = \tilde{N}_q)$, across subjects, estimated for models M_q ($q = 1$ to 8, the number of models, including the reference model). The best-performing estimation models for stiffness were models based upon detecting peak force during the probing motion and dividing this peak value by the amount of penetration into the boundary, as well as models based on local linear regression of force and displacement when the data in the inward portion of the probing motion are considered, i.e. the data from the point of penetration to the surface up until maximum displacement occurs (the gray circle on figure 4b). The predictions of Model #5 are shown in Figure 8, the high correspondence to the subject’s performance in Figure 6 is evident.

A very different result was obtained using data from the second, outward going, portion of the movements. In this case the probability is quite low. A remarkable finding was that using the complete history of force and position and performing a linear regression did not lead to good agreement with the subjects’ responses, as shown by the poor score of the “Global Stiffness”. Poor results were generally obtained when using a symmetric data set around the point of maximum penetration. We conclude from this analysis that subjects were likely to perform an estimate of stiffness based on a small subset of the experienced force/motion information.

Models were ranked by the fraction of trials in which model’s M_q prediction agreed with the subject response, $P(N = \tilde{N}_q)$. A generalized linear mixed model (McCul-




Fig. 8. The psychometric curves from Peak F/P_{per} model (# 5), is shown in this figure. The legend displays the different delays. The same tendency as the sigmoid curves created for subject are evident.

loch and Searle, 2001; Fitzmaurice, Laird, and Ware, 2004) was used to estimate the probability of agreement and associated standard errors across subjects for each level of delay, while accounting for the correlation between repeated measurements.

3.3. Penetration changes with delay

The mean nominal penetration and effective penetration for all subjects is presented in Figure 9. The small x 's are the median penetration for each trial and the larger symbol is the median value for each delay. It is obvious that the amount of nominal penetration increases with delay (top graph). This implies that the maximum force delivered by the surface to the subject, on average, was highest in the 60 ms delay condition and smallest for the -50 ms delay condition. This finding excludes the possibility that trajectories at different delays have the same profile, on average, over the force/displacement plane. The results of Figure 9 show that, with negative delays, the nominal penetration is smaller than with zero delay ($P < 0.001$, $F = 474.15$). Conversely, with positive delays nominal penetration is higher and effective penetration is smaller. Accordingly, the values for maximum force (the position of the black circle in Figure 4b) follow the same trend. As stiffness is the ratio of force to displacement, the increase in the experienced peak force would tend to increase the perceived level of stiffness. This effect may be enhanced by a trend toward reduced effective penetration with increasing delay ($P < 0.001$, $F = 64.88$), as smaller penetration, for a given force, would correspond to higher stiffness. The aver-




Fig. 9. The amount of nominal penetration (top) and effective penetration (bottom) for all subjects. The large marking represent the median value across all subjects, whereas each point in the black bars is a value in a single trial.

age duration of a single back and forth hand movement within the object's boundary was 430 ± 130 ms. The corresponding frequency was 2.2 Hz. The stiffness of the surface, which determines the amount of interaction force, and the delay incorporated into the surface didn't have a significant effect on the movement duration.

3.4. An interaction with asymmetric delay excluded additional three models (#2 Max Force, #4 F/P model, and #7 Local stiffness on the way in) leaving model #5 Peak F/P_{per} as the best candidate to explain the subjects behavior

The interactions with simple linear stiffness boundaries allow us to exclude models whose predictions are incompatible with the subjects' responses. In particular, estimates based on linear regressions over the entire set of data and estimates based on the amount of penetration generated unacceptable predictions. The "surviving models" are:

- Maximum force (Model #2).
- Peak force/penetration ratio (Peak F/P , Model #4).
- Peak force/penetration ratio with boundary estimation (Peak F/P_{per} , Model #5).
- Local stiffness on the way in (Model #7).

To further reduce this list, we tested subjects in a second experiment, involving the interaction with a more complex

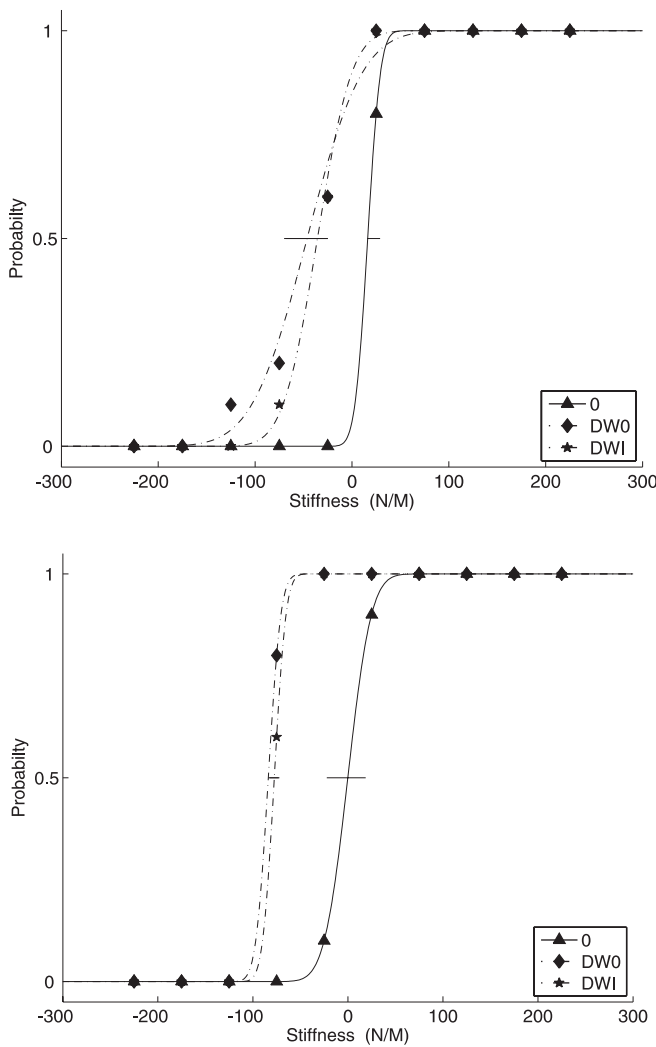


Fig. 10. Psychometric curves derived for one of the subjects for the three conditions (Top). The output of the Peak F/P_{per} , model #5, for the same subject (bottom). It is clear that the model predicts the tendency of the perceived shift in stiffness.

boundary. In this experiment (see Methods, subsection 2.1) the stiffness of the boundary was affected by the delay in different ways, depending on the direction of hand motion. The delay was either present on the inward probing motion and not in the outward motion (“delayed-way-in”), or vice versa – the inward motion was not delayed and the outward motion was delayed (“delayed-way-out”).

The shift in perception for all five subjects was toward overestimation of the stiffness for both cases, delayed-way-in and delayed-way-out. The change in perception, as evidenced by the shift of the psychometric curves of one of the subjects (top), along with the prediction of model Peak F/P_{per} , #5 bottom are shown in Figure 10.

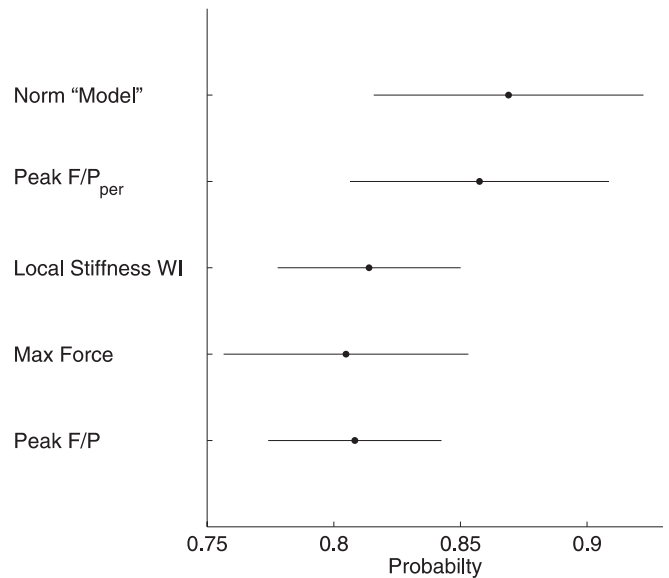


Fig. 11. The probability of agreement for the various models over all three conditions (no delay, delay-on-the-way-in and delay-on-the-way-out) of the second experiment. The black dots are the mean value across subjects and the wings represent 95% confidence intervals. The WI on the Local Stiffness estimators line stand for way in.

The probability of agreement for the four models and the reference model is displayed in Figure 11. The highest value is for model Peak F/P_{per} , model #5, (and is the closest to the Reference Model). The low value given by Peak F/P , model #4 (although it is strongly resembles model #5) is probably due to the tendency of the model to overestimate the stiffness in the delayed-way-in condition. It is worth mentioning that, for the maximum force model, the shift in perception on the delayed-way-in condition is due to the experience of higher peak forces ($T - test, P < 0.001$).

4. Discussion

We have investigated the effect of feedback delays on the perception of stiffness (Jones and Hunter, 1990). Our paradigm is consistent with the general setting of bilateral teleoperation, in which an operator moves the effector of a manipulator and receives the force reflected by an object touched by a remotely controlled system (Niemeyer and Slotine, 1991; Anderson and Spong, 1989).

We carried out two experiments to assess the effect of delay on the perception of stiffness. In these experiments, subjects interacted mechanically with a virtual boundary, with programmable elastic properties. Subjects held a planar manipulator, and the force at the interface between the hand and the environment was programmed to emulate different delay

patterns with respect to the hand position. The first experiment introduced a constant lag or lead between interface force and hand position. In the second experiment, the delay was introduced either as the hand moved into the elastic surface or out, but not in both directions.

In the first experiment, we found that subjects tended to overestimate stiffness when the force lagged the hand motion and to underestimate stiffness when the force led the hand motion. This result is limited to the tested conditions, i.e. a spring-like surface with a small range of delays, between -50 and $+60$ ms. Larger delays tended to disrupt both the performance and the ability of subjects to discriminate stiffness.

We investigated the behavior of various estimation models, and compared the predicted behavior to the subjects' answers. Here, we considered a set of very specific models, each one having the property of producing the correct stiffness comparisons in the absence of delays. Surprisingly, poor results were observed when the entire history of the force and hand position data was used for linear regression. Comparison of model outputs and the subjects' responses favors models based on a limited set of force and position data. Specifically, we found that subjects' responses could be reasonably reproduced by assuming that stiffness is estimated by one of the following schemes: (a) dividing the maximum experienced reaction forces by the perceived penetration inside the boundary; (b) reporting the actual maximum force as a proxy for stiffness; or (c) on the basis of a local regression of force and position over data collected during the initial inward probing motions.

In the second experiment, we found a consistent tendency for all subjects to overestimate the stiffness when the interaction force was delayed either to the inward or to the outward hand motion. This finding excludes models based only on data from the way into the surface or on the maximum force as proxy for stiffness and leave only two models, both dividing the maximum perceived force by the amount of penetration (#4 and #5), which had the best average probability of agreement with the subjects (see Figure 11).

While our best fitting models may reflect the operation of the haptic perceptual system, there are clearly many other possible models and estimation methods that we didn't consider. As is often the case with experimental studies, the most conclusive results of this one concern the models that failed to reproduce the data, and that therefore can be ruled out as biologically plausible models.

The best model was the one dividing the maximum interaction force by an estimate of the perceived penetration. The discrepancy between the effective penetration and the perceived penetration can be regarded as an attraction toward a prior representation of the boundary. During more than half the probing motions, the boundary was in a static position (this was the case with zero delay between force and position). This model assumes that the haptic system holds a representation of the boundary's position. When probing the surfaces in which a delay between force and position is present, the boundary

shifts towards (or away from) the subject (depending on the sign of τ , see equation 1). Then, the model suggests that when the subject needs to assess the amount of penetration, the internal representation of the boundary biases the measurement. The combination of prior information and sensed data by the brain has recently been addressed by Koerding et al. (Koerding and Wolpert, 2004; Koerding, pi Ku, and Wolpert, 2004) in the context of position and force representation within a Bayesian framework. Miyazaki et al. (Miyazaki, Nozaki, and Nakajima, 2005) considered the same framework in the representation of timing. Further studies are required to explore the construction of this internal representation.

Perception of stiffness with delay could be considered as a problem of "binding" force and position and of simultaneity in perception (Stone et al., 2001). Simultaneity is the occurrence of two or more events at the same time. Since the nervous system is likely to process each stimulus modality at different times, some mechanism is probably responsible for perceptually co-aligning these modalities and creating a coherent picture of the present. Various studies have been made of cross-modal simultaneity. Among them are binding of the audio and visual systems (Shams, Kamitani, and Shimojo, 2002; Calvert, 2001; Fujisaki et al., 2004) visual and haptic (Violyentev, Shimojo, and Shams, 2005; Ernst and Banks, 2002), and audio and haptic (Adelstein et al., 2003), (Levitin et al., 1999). This work was aimed at exploring the interactions between the modalities within the haptic and motor systems. Specifically, the positioning of the limb along the trajectory was considered as one modality and the estimation of interaction force was considered as the other. During interaction with a surface in which the force and position are not time shifted, the instants of maximum force and maximum penetration are coincident. It would be plausible to assume that the haptic system estimates the stiffness of a boundary by relating the experienced peak force with either the preprogrammed extent of penetration in the surface or with the actual maximum penetration. However, neither of these possibilities is supported by our analysis of the data. The model which best captured the subjects' responses, suggests that, along with the estimation of peak force, an estimate of the amount of penetration up to that point is carried out. It has been shown that self-generated action might lead to attenuation of sensation from the moving part (Chapman et al., 1987). Voss et al (Voss et al., 2006) used transcranial magnetic stimulation to study this effect and concluded that this phenomenon might be attributed to an earlier planning mechanism in the motor cortex. This fact could definitely influence the perception of stiffness. However, it does not seem to be sufficient to explain our results, in which both overestimation and underestimation of stiffness are observed for positive and negative delays. During the experiment, a limited number of delays between the force and position were tested. The delay causes a discrepancy between the two points of maximum force and maximum penetration. We observed that positive delays (force lagging position) greater than 80 ms lead to an abol-

ishment of the perception of a surface. The actual two values which were tested were 80 and 110 ms. These are relatively large compared to the duration of the complete in-out movement 430 ± 130 ms. We believe that delays that are a major part of the whole motion cause too big a mismatch between position, movement direction and the force. This mismatch leads to an abolishment of the perception of a surface. This suggests that 60 ms may define a temporal window within which the two sensory events (detection of peak force and of peak penetration) can be considered as simultaneous. The temporal extent of simultaneity defines the effective temporal extent of what is commonly called the present time or “now”. In abstract terms, “now” is a zero-measure boundary between past and the future. However in the perceptual system this separation has a non-zero width, and the temporal window established by our experiments provides a first estimate of this width (Stone et al., 2001). Though these delays are rather small compared to the demands imposed by large scale teleoperation systems, we would like to stress their importance. First, a delay of 50 ms is equivalent to signals traveling at the speed of light for a distance of several thousands kilometers. Such scenarios are likely to occur in future telesurgery or remote operation in hazardous areas. Second, ways to compensate for delays such as wave variables (Niemeyer and Slotine, 1991), Smith predictor (Miall et al., 1993) and Kalman filters (Kalman and Bucy, 1960) need to be found. It can be assumed that the delay canceling would not be perfect and a mismatch error (between the actual delay and the one predicted by the cancellation mechanism) would occur. If the error is of the order of tens of milliseconds, the results of this work might explain how stiffness perception would be affected (Foulkes and Miall, 2000; Fujisaki et al., 2004; Stone et al., 2001).

One should note that these compensation methods have their limitations and advantages: wave variables (Niemeyer and Slotine, 1991) stabilize delayed systems at the expenses of performance, while predictive techniques such as Smith predictor (Miall et al., 1993) can compensate for delay but might be unstable if the prediction is wrong. Further studies are required to explore the effect of these techniques on the perception of stiffness.

This study could be extended in many directions. An interesting open question concerns the ability of the brain to adapt to delays between position and force. The ability to perceive and adapt to visual feedback delays was studied by a few researchers (Foulkes and Miall, 2000; Fujisaki et al., 2004; Stone et al., 2001), however, the study of adaptation to delay in haptic interfaces has only recently begun (Vogels, 2004). One should also note that haptics involves tactual sensation, proprioception, and the distinction between active and passive touch (Srinivasan and LaMotte, 1995). All these haptics components should be taken into account in a complete study of perception of and adaptation to delays, which is of critical importance for developing telerobotics, telesurgery and telepresence technologies.

Acknowledgements

This research was supported by Grant No. 2003021 from the United States-Israel Binational Science Foundation (BSF), Jerusalem, Israel and by NINDS grant NS35673.

References

- Adams, R. J. and Hannaford, B. (2002). Control law design for haptic interfaces to virtualreality. *Control Systems Technology, IEEE Transactions on* **10**(1): 3–13.
- Adelstein, B. D., Begault, D. R., Anderson, M. R., and Wenzel E. M. (2003). Sensitivity to haptic-audio asynchrony. In *Proceedings, 5th International Conference on Multimodal Interfaces, ACM.*, pp. 73–76, Vancouver, Canada.
- Anderson, R. J. and Spong, M. W. (1989). Bilateral control of teleoperators with time delay. *Automatic Control, IEEE Transactions on* **34**(5): 494–501.
- Calvert, G. A. (2001). Crossmodal processing in the human brain: insights from functional neuroimaging studies. *Cereb Cortex* **11**(12): 1110–1123.
- Chapman, C. E., Bushnell, M. C., Miron, D., Duncan, G. H., and Lund, J. P. (1987). Sensory perception during movement in man. *Exp Brain Res* **68**(3): 516–524.
- Chatterjee, S. and Hadi, A. S. (1986). *Influential Observations, High Leverage Points, and Outliers in Linear Regression*. Statistical Science.
- Conditt, M. A. and Mussa-Ivaldi, F. A. (1999). Central representation of time during motor learning. *Proc Natl Acad Sci USA* **96**(20): 11625–11630.
- Ernst, M. O. and Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature* **415**(6870): 429–433.
- Fitzmaurice, G. M., Laird, N. M. and Ware, J. M. (2004). *Applied Longitudinal Analysis*. Wiley, New York.
- Flash, T. and Gurevich, I. (1997). *Self-Organization, Computational Maps, and Motor Control*, chapter Models of motor adaptation and impedance control in human arm movements, pp. 423–481. Elsevier Science.
- Foulkes, A. J. and Miall, R. C. (2000). Adaptation to visual feedback delays in a human manual tracking task. *Exp Brain Res* **131**(1): 101–110.
- Fujisaki, W., Shimojo, S., Kashino, M. and Nishida, S. (2004). Recalibration of audiovisual simultaneity. *Nat Neurosci* **7**(7): 773–778.
- Golub, G. H. and Van Loan, C. F. (1996). *Matrix Computations*. Johns Hopkins Studies in Mathematical Sciences.
- Grewal, M. S. and Andrews, A. P. (2001). *Kalman Filtering: Theory and Practice Using MATLAB, 2nd ed.* Wiley-Interscience.
- Hannaford, B. (1989). A design framework for teleoperators with kinesthetic feedback. *Robotics and Automation, IEEE Transactions on* **5**(4): 426–434.

- Hirche, S., Bauer, A., and Buss, M. (2005). Transparency of haptic telepresence systems with constant time delay. In *Control Applications, 2005. Proceedings of 2005 IEEE Conference on*, pp. 328–333.
- Jones, L. A. and Hunter, I. W. (1990). A perceptual analysis of stiffness. *Exp Brain Res* **79**(1): 150–156.
- Kalman, R. E. and Bucy, R. S. (1960). New results in linear filtering and prediction theory. *Journal of Basic Engineering ASME Transactions* **83**: 95–107.
- Karniel, A. and Mussa-Ivaldi, F. A. (2003). Sequence, time, or state representation: how does the motor control system adapt to variable environments? *Biol Cybern* **89**(1): 10–21.
- Koerding, K. P., Ku, S. and Wolpert, D. M. (2004). Bayesian integration in force estimation. *J Neurophysiol* **92**(5): 3161–3165.
- Koerding, K. P. and Wolpert, D. M. (2004). Bayesian integration in sensorimotor learning. *Nature* **427**(6971): 244–247.
- Lackner, J. R. and Dizio, P. (1994). Rapid adaptation to Coriolis force perturbations of arm trajectory. *J Neurophysiol* **72**(1): 299–313.
- Levitin, D. J., MacLean, K., Mathews, M., Chu, L., and Jensen, E. (1999). The perception of cross-modal simultaneity. *International Journal of Computing Anticipatory Systems*.
- Libet, B., Wright, E. W., Feinstein, B. and Pearl, D. K. (1979). Subjective referral of the timing for a conscious sensory experience: a functional role for the somatosensory specific projection system in man. *Brain* **102**(1): 193–224.
- Massaquoi, S. G. and Slotine, J.-J. E. (1996). The intermediate cerebellum may function as a wave-variable processor. *Neurosci Lett* **215**(1): 60–64.
- McCulloch, C. E., and Searle, S. R. (2001). *Generalized, Linear, and Mixed Models*. Wiley, New York.
- Miall, R., Weir, D., Wolpert, D. M. and Stein, J. (1993). Is the Cerebellum a Smith Predictor? *J Mot Behav* **25**(3): 203–216.
- Miyazaki, M., Nozaki, D. M. and Nakajima, Y. (2005). Testing Bayesian models of human coincidence timing. *J Neurophysiol* **94**(1): 395–399.
- Mussa-Ivaldi, F. A., Hogan, N. and Bizzi, E. (1985). Neural, mechanical, and geometric factors subserving arm posture in humans. *J Neurosci* **5**(10): 2732–2743.
- Niemeyer, G. and Slotine, J.-J. E. (1991). Stable adaptive teleoperation. *IEEE Journal of Oceanic Engineering* **16**(1): 152–162.
- Niemeyer, G. and Slotine, J.-J. E. (2004). Telemanipulation with Time Delays. *The International Journal of Robotics Research* **23**(9): 873–890.
- Shadmehr, R. and Mussa-Ivaldi, F. A. (1994). Adaptive representation of dynamics during learning of a motor task. *J Neurosci* **14**(5 Pt 2): 3208–3224.
- Shams, L., Kamitani, Y. and Shimojo, S. (2002). Visual illusion induced by sound. *Brain Res Cogn Brain Res* **14**(1): 147–152.
- Srinivasan, M. A. and LaMotte, R. H. (1995). Tactile discrimination of softness. *J Neurophysiol* **73**(1): 88–101.
- Stone, J. V., Hunkin, N. M., Porrill, J., Wood, R., Keeler, V., Beanland, M., Port, M. and Porter, N. R. (2001). When is now? Perception of simultaneity. *Proc Biol Sci* **268**(1462): 31–38.
- Violentyev, A., Shimojo, S. and Shams, L. (2005). Touch-induced visual illusion. *Neuroreport* **16**(10): 1107–1110. A talk about neurological pathways between the audio and visual cortex.
- Vogels, I. M. L. C. (2004). Detection of temporal delays in visual-haptic interfaces. *Hum Factors* **46**(1): 118–134.
- Voss, M., Ingram, J. N., Haggard, P. and Wolpert, D. M. (2006). Sensorimotor attenuation by central motor command signals in the absence of movement. *Nat Neurosci* **9**(1): 26–27.
- Wichmann, F. A. and Hill, N. J. (2001a). The psychometric function: I. Fitting, sampling, and goodness of fit. *Percept Psychophys* **63**(8): 1293–1313.
- Wichmann, F. A. and Hill, N. J. (2001b). The psychometric function: II. Bootstrap-based confidence intervals and sampling. *Percept Psychophys* **63**(8): 1314–1329.