Behavioral/Systems/Cognitive

How Soft Is That Pillow? The Perceptual Localization of the Hand and the Haptic Assessment of Contact Rigidity

Assaf Pressman,[1,2] **Amir Karniel,**[2] **and Ferdinando A. Mussa-Ivaldi**[1,3]

[1]Sensory Motor Performance Program, Rehabilitation Institute of Chicago, Chicago, Illinois 60611, [2]Department of Biomedical Engineering, Ben-Gurion University of the Negev, Beer-Sheva, 84105 Israel, and [3]Department of Physiology, Northwestern University Feinberg School of Medicine, Chicago, Illinois 60611

A new haptic illusion is described, in which the location of the mobile object affects the perception of its rigidity. There is theoretical and experimental support for the notion that limb position sense results from the brain combining ongoing sensory information with expectations arising from prior experience. How does this probabilistic state information affect one's tactile perception of the environment mechanics? In a simple estimation process, human subjects were asked to report the relative rigidity of two simulated virtual objects. One of the objects remained fixed in space and had various coefficients of stiffness. The other virtual object had constant stiffness but moved with respect to the subjects. Earlier work suggested that the perception of an object's rigidity is consistent with a process of regression between the contact force and the perceived amount of penetration inside the object's boundary. The amount of penetration perceived by the subject was affected by varying the position of the object. This, in turn, had a predictable effect on the perceived rigidity of the contact. Subjects' reports on the relative rigidity of the object are best accounted for by a probabilistic model in which the perceived boundary of the object is estimated based on its current location and on past observations. Therefore, the perception of contact rigidity is accounted for by a stochastic process of state estimation underlying proprioceptive localization of the hand.

Introduction

An illusion is a distortion of the senses, revealing how the brain organizes and interprets sensory stimuli. Although illusions are subjective, they are generally shared by most people (Solso, 2001; Levitin, 2002). Illusions can be referred to a single modality, e.g., Kanizsa triangle (Kanizsa, 1955; Renier and Volder, 2005), or they may span multiple modalities and involve timing differences between them, such as in the McGurk effect (McGurk and MacDonald, 1976; Munhall et al., 1996) or the rubber hand illusion (Botvinick and Cohen, 1998; Peled et al., 2003). Various illusions are known to appear within the sensory-motor system such as the phantom limb (Halligan, 2002) and the cutaneous rabbit illusion (Geldard and Sherrick, 1972). Here, we present a new illusion involving the motor system. In this illusion, the perception of an object's rigidity or stiffness is distorted. The perception of rigidity is very important in determining how we manipulate things in our environment. For example, we apply different contact forces to a glass versus a Styrofoam cup.

A variety of works in the field of haptics have demonstrated the interaction between force and position information in stiffness assessment within the nervous system (Jones and Hunter, 1990; Srinivasan and LaMotte, 1995; Pressman et al., 2007). Additionally, a number of illusions within this field have been reported (Jones, 1988) and, more recently, neural correlates to the size–weight illusion and the material–weight illusion (Chouinard et al., 2009) have been found. Here, we explore how the brain estimates the stiffness of a contact by combining past and current position information.

We applied an unexpected displacement of the object boundary and evaluated its effect on the perception of rigidity. Based on our previous studies (Pressman et al., 2007), we assumed that after a prolonged interaction with a stationary object boundary, a subject's nervous system forms a representation of that boundary's location. This representation would then bias the perception of the object's position with respect to the subject. Following this view, a small displacement of the boundary would cause a change in the perceived rigidity of the object because of the change in perceived distance traveled inside the boundary. The displacement would therefore create an illusion of a change in stiffness, when actually no such change took place.

We found that subjects did indeed experience the illusion. Then we considered the possibility that the brain derives the location of the boundary by combining the ongoing sensory information about the contact with a prior belief, derived from past experience. We designed a second experiment in which the object was smoothly shifted back and forth. Trial-by-trial learning in the context of random force perturbations was found to be well represented by an autoregressive process (Scheidt et al., 2001), which combines current input and performance with past performance.

Together, our findings demonstrate for the first time that the same adaptive mechanisms that are used by the brain to optimally

Received Sept. 1, 2010; revised Feb. 18, 2011; accepted March 1, 2011.

Author contributions: A.P., A.K., and F.A.M.-I. designed research; A.P. performed research; A.P., A.K., and F.A.M.-I. analyzed data; A.P., A.K., and F.A.M.-I. wrote the paper.

This work was supported by National Institute of Neurological Disorders and Stroke Grant 2R01NS035673-10A2 and by Grant 2007195 from the United States–Israel Binational Science Foundation.

Correspondence should be addressed to Assaf Pressman, Biomedical Engineering Department, Ben Gurion University of the Negev, Beer Sheva, 84105 Israel. E-mail: pressman@bgu.ac.il.

DOI:10.1523/JNEUROSCI.4656-10.2011

 0270-6474/11/316595-10$15.00/0

estimate the state of motion of the arm based on combining current information and prior knowledge may also have a direct impact on the perception of the mechanical properties of the environment we come into contact with. Preliminary results regarding this illusion were presented in Pressman et al. (2008).

Materials and Methods

Setup

Thirteen human subjects (eight males; five females; age, 27 ± 7 years) participated in this experiment after signing the informed consent form approved by Northwestern's Institutional Review Board. Seven subjects participated in the first experiment (constant reference experiment) and six subjects participated in the second experiment (moving reference experiment) (Table 1). Subjects held, with their dominant (right) hand, the handle of a two-degrees-of-freedom robotic manipulandum, and looked at a screen, placed horizontally above their hand (for further details about the robotic manipulandum, see Shadmehr and Mussa-Ivaldi, 1994; Conditt et al., 1997).

We used a forced-choice paradigm to assess the effect of boundary displacement on the perception of stiffness. In each trial, two surfaces were rendered by a robotic device (Fig. 1*A*), sequentially, one at a time. Force was exerted perpendicular to the coronal (frontal) plane, with the subject facing the robotic device. Subjects were asked to alternate contacts with both surfaces by moving the handle of the manipulator against their boundaries. Movements were made in a sagittal plane passing through the subject's shoulder and perpendicular to the transverse (horizontal) plane (motion was limited to the transverse plane due to the mechanical design of the robotic device). We call these testing contacts "probing motions." After performing as many probing motions as they wished, subjects were to report which of the two surfaces they perceived to be stiffer. Subjects were not instructed about extent of penetration inside the boundaries.

The elastic force generated by the manipulandum was proportional to the distance of the subject's hand from the surface boundary. On average, the boundary of the object was 45 cm from the subject's shoulder, along the sagittal plane. Occasionally, one of the two surfaces (D surface) was slightly displaced toward or away from the subject. The other surface (K surface) was kept stationary across trials. The subjects did not see their hands. A line perpendicular to the object's border was presented to indicate the lateral location of the hand, but no visual information was available to the subject about the vertical distance of the hand to the border. A circular disk (14 cm diameter) was projected on the screen in alternating red and green colors (Fig. 1*B*). The colors provided the subjects with an identifier for two alternating virtual surfaces, referred to as "red surface" and "green surface." The shape and location of the disks were chosen not to disclose any spatial information regarding the surfaces. The two colors, however, were assigned randomly in each trial, so that each surface type (K or D) was not uniquely associated with a color.

Subjects probed each of the two surfaces several times. They switched between the two until they felt ready to give their response to the question "Which surface is stiffer (green or red)?" by pressing one of two buttons on a handheld device. A single probing sequence is shown in Figure 1*C*. Switching took place once the subject had retracted his/her hand >8 cm from the boundary (Fig. 1*C*, red squares).

We acquired three types of data, as follows: (1) the responses of the subjects (green or red, for each trial); (2) the position of the hand along the *y*-axis, sampled at a rate of 100 samples per second. Position data were used in real time to generate the virtual surfaces; and (3) the interaction force with the surface. This force was calculated in real time based on the hand position, the displacement and the elastic properties of the surface (stiffness and boundary).

Table 1. Log likelihood values for the CP and the VP model

Model subject number	Experiment	Displacement direction	Likelihood value CP model	VP model	VP − CP	Mean VP − CP
1	Constant reference	+	46.68	44.56	−2.13	−6.01
2		+	34.47	34.01	−0.480	
3		+	39.76	21.67	−18.09	
4		−	82.75	88.94	6.20	
5		−	65.48	61.77	−3.72	
6		−	63.10	56.75	−6.35	
7		+	48.12	30.55	−17.56	
8	Moving reference	+	80.43	51.37	−29.07	−17.72
9		+	36.95	32.50	−4.45	
10		+	76.31	65.60	−10.71	
11		+	62.67	37.54	−25.14	
12		+	64.59	46.22	−18.37	
13		+	52.98	34.10	−18.89	

Displacement direction is given for each subject as + (toward) or − (away) signs. VP − CP is the difference between VP and CP. Mean VP − CP shows the mean value across VP − CP for all subjects in each experiment.

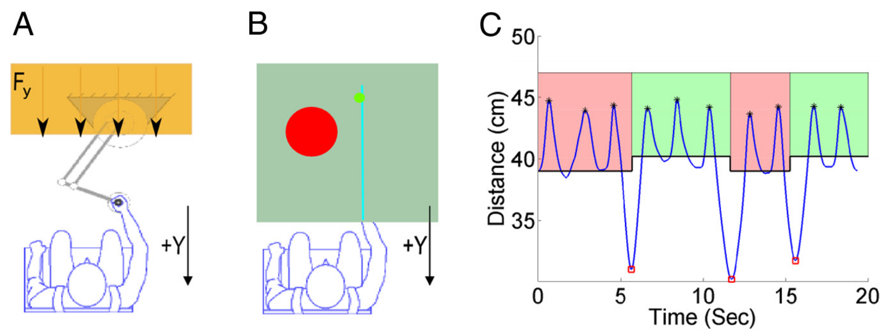

Figure 1. ***A***, Subject holding the handle of the manipulandum and interacting with the object. Object (orange rectangle) is oriented parallel to the coronal plane and located ~40 cm away from the subject. Direction of force is toward the subject and it is perpendicular to the coronal plane (arrows). Displacement of the object would keep it parallel to the coronal plane and would either move it toward or away from the subject. Subjects were advised to probe the object by moving the hand within a plane parallel to the sagittal plane and passing through their right shoulder. The motion was free, that is, no mechanical limits were deployed to keep the subjects on this plane. ***B***, Visualized environment during the experiment. The rectangle represents the screen on which information was displayed. This screen blocked any visualization of the arm. The red circle represents the surface with which the subject is currently interacting. The cyan line represents the lateral location of the arm and serves as the guideline to arm placement. The green dot marks the advised arm configuration during probing. ***C***, A typical probing of the boundary. Rectangles show the different surfaces (pink, K; green, D). Note that the green D surface is displaced away from the subject in this case. The black line shows the location of the boundary for each surface. Asterisks on the maximum points represent a single probing into the surface. Red squares mark the swapping of the two surfaces as subjects retracted their arm by >8 cm away from the surface's boundary.

Experiment 1: constant reference. The force exerted by the virtual surfaces was, in the *y*-axis direction, in proportion to the displacement from the boundary, $Y_0{}^n$, during the current (n) trial.

$$F_y = \begin{cases} -k[Y - (Y_0{}^n + H)] & Y < Y_0{}^n + H \\ 0 & Y \geq Y_0{}^n + H \end{cases}, \tag{1}$$

where F_y is the force in the y-axis direction, k is the spring's stiffness constant, Y is the position along the y-axis, and $Y_0{}^n$ is the coordinate of the boundary. During the first experiment, $Y_0{}^n = 0$ for all n.

During each trial 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 were never equal. The D surface was displaced toward or away from the subject in a portion of the trials, which was chosen randomly, during the course of the experiment and was the same for all subjects. The displacement was implemented by changing the value of H in Equation 1 for the whole extent of the specific trial. Values for H were changed from zero to a value drawn from a Gaussian distri-

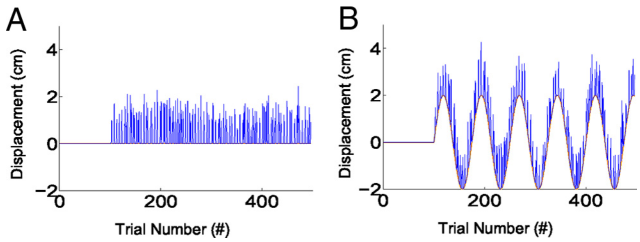

Figure 2. ***A***, ***B***, Baseline and displacements in both boundaries, K (orange line) and D (blue line). ***A***, Setup for the constant reference experiment, where the baseline position of both the boundaries was not displaced. ***B***, A sinusoidal displacement of both boundaries is evident starting on trial 100.

bution ($\mu = \pm 1.5$ cm, $\sigma = 0.35$ cm) only for the D surface. Each subject (three toward, four away in the constant reference experiment; six toward, none away in moving reference experiment) encountered either a positive or a negative displacement (μ). The K surface was set to $Y_0{}^n$ and was never displaced.

Experiment 1: constant reference consisted of two blocks. The first was a reference block in which no displacement of the D boundary took place and lasted for 100 trials. During the second block (400 trials), on random trials spaced two to six trials apart, the D boundary was displaced toward or away from the subject (for a specific subject, the direction of displacement was kept fixed during the whole experiment) (Fig. 2*A*).

Experiment 2: moving reference. In the second experiment (with six subjects), an additional slow sinusoidal displacement was applied to the position of both the K and D surfaces; specifically, the value of $Y_0{}^n$ in Equation 1 was not constant across the whole experiment but changed as a function of the trial number n as follows (units are in centimeters) (Fig. 2*B*):

$$Y_0{}^n = \begin{cases} 0 & n < 100 \\ \sin(2 \cdot \pi \cdot (n - 100)/75) \cdot 2 & n \geq 100 \end{cases}. \quad (2)$$

The slow trial-by-trial sinusoidal displacement applied to the objects in this experiment resulted in a maximum step change (relative to the previous trial) of 0.17 cm. This was small compared with the mean of the distribution from which standard object displacements were chosen (1.5 cm). However, the peak-to-peak amplitude of the slow sinusoidal displacements was 4 cm and thus exceeded the mean standard object displacement. As shown below in Results, these conditions allowed us to distinguish between a model which can account for slow changes in surface location and one which cannot.

Psychometric curves

Figure 3*C* schematically illustrates psychometric curves describing subject replies as function of the gap between the stiffness of the surfaces D and K. These curves estimate the frequency with which subjects indicated that the D surface was stiffer than the K surface, as a function of the actual stiffness difference, $K_D - K_K$. For two linear spring-like surfaces without a displacement (i.e., the value of H in Equation 1 for the D surface is equal zero), we expected the answers to reach chance (value of 0.5) level when the two values of the stiffness were equal. This is the point of subjective equality (PSE) (Gescheider, 1997; Ohnishi and Mochizuki, 2007). A sigmoid-shaped curve is expected, as answers are likely to be correct when the stiffness levels are substantially different. If a change is made to one of the surfaces that alters the perception of stiffness, the psychometric curve would not be centered on $K_D - K_K = 0$ anymore, but would shift to either side, depending on the perception being biased toward stiffer or more compliant evaluation.

Specifically, we let $\Delta = K_D - K_K$ be the difference between the D and K stiffness values. Each point on the psychometric curve is an estimate of the probability of the subject reporting that the D surface is stiffer than the K surface as a function of Δ.

$$\hat{P}_S(\text{Reporting that } \hat{K}_D > \hat{K}_K | \Delta) = \frac{\sum_{j=1}^{R(\Delta)} N^j}{R(\Delta)}, \quad (3)$$

where

$$N^j = \begin{cases} 1 & \text{reply:} \quad \hat{K}_D^j > \hat{K}_K^j \\ 0 & \text{reply:} \quad \hat{K}_D^j < \hat{K}_K^j \end{cases}, \quad (4)$$

and 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 current block. The caret over the stiffness of the K and D surfaces stands for the stiffness perceived by the subject, as opposed to the actual nominal stiffness of the surface. The subindex S in $\hat{P}_S$ designates this function for subject S. As described below, we derived similar functions for each computational model.

The psychometric curves were estimated separately for three blocks for each subject on both experiments. The first block was for reference, the second block was either a trial with no displacement at the D surface or a trial with displacement at the D surface.

The stiffness at the PSE (where $\hat{P} \approx 0.5$) was evaluated using a maximum likelihood fit of a sigmoid function to the data points. The bootstrap method was used to estimate the goodness of fit (Wichmann and Hill, 2001a,b).

Logistic regression

As explained in the previous section, subjects became more uncertain in assessing the relative rigidity of the surfaces when their stiffness became more alike (Fig. 3C, shading). The logit transformation (Hosmer and Lemeshow, 2000) provides us with a way to describe the probability of the subject reporting that the D surface is stiffer than the K surface, as a continuous function of the true stiffness difference:

$$P(\Delta) = \frac{1}{1 + e^{-g(\Delta,\bar{\beta})}}. \quad (5)$$

Here, we assumed that:

$$g(\Delta,\bar{\beta}) = \bar{\beta}^{\mathrm{T}} \cdot \bar{\Delta}; \; \bar{\beta} = [\beta_1,\beta_2]^{\mathrm{T}}; \; \bar{\Delta} = [1,\Delta]^{\mathrm{T}}. \quad (6)$$

Generally, the function $P(\Delta)$ has the shape of a sigmoid, approaching zero for extremely negative values of Δ and approaching one for extremely positive values. The transition between the two values takes place in the intermediate region, whose size and location depends on the vector $\bar{\beta}$.

We estimated the parameter $\bar{\beta}$ using the logistic regression procedure. There, the dichotomous response ($r \in \{0,1\}$) by the subject regarding the difference in stiffness between the two surfaces was regressed with the continuous nominal stiffness differences between the two surfaces across all trials (Δ) (Hosmer and Lemeshow, 2000). For each given dataset, we used a nonlinear least-squares estimate to assess the values of the parameter. Once these values had been estimated, the probability that the given dataset was consistent with the estimated parameters (or the likelihood) was estimated by the binomial likelihood function:

$$l(\bar{\beta}) = \prod_{i=1}^{n} p(\Delta_i)^{r_i}[1 - p(\Delta_i)]^{1-r_i}. \quad (7)$$

Since the likelihood function is monotonically increasing and usually takes very small values, it is more convenient to consider $L = -2\ln(l)$. Smaller values of L correspond to a better fit or a better description of the data by the parameters.

The vector $\bar{\beta}$ offers a compact description of the dataset. As shown above, this vector describes a sigmoidal curve. We used the likelihood value to compare the performances of several models. To this end, we estimated the likelihood to observe a response r, assuming that a model ($\bar{\beta}$) is true and Δ is known: $P(r|\bar{\beta}, \Delta)$. We identify the model that achieves the highest likelihood (lowest negative log of likelihood) as the best model (MacKay, 2002).

Prediction

Based on previous studies we employ the following model for the perception of stiffness (Pressman et al., 2007).

$$K = \frac{\max(\mathrm{Force})}{B_{\mathrm{perceived}}}. \tag{8}$$

The maximum force encountered during the whole probing motion was normalized by the perceived amount of penetration into the surface ($B_{\mathrm{perceived}}$). This amount was not necessarily equal to the actual distance traveled from the surface boundary. A change in the perceived amount of penetration might alter the perception of stiffness.

We assumed that during numerous interactions with a surface at the same fixed position, an internal representation of the surface's boundary was created within the nervous system. This might have happened during the first block of the experiment, where no displacements took place. Such representation would alter the perceived penetration when the boundary of the surface was displaced away from that previous zero point. This in turn would influence the estimation of stiffness.

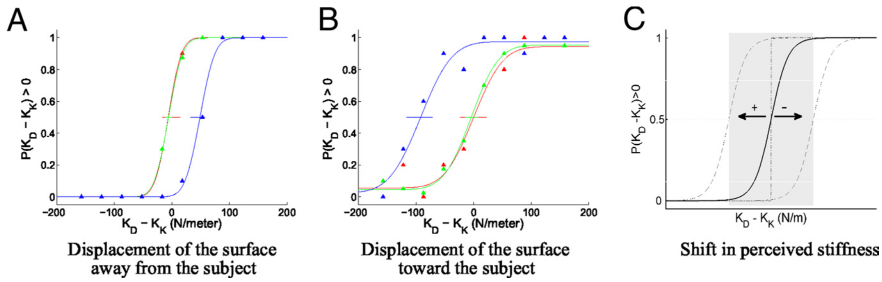

Figure 3. ***A***, ***B***, The estimated probability that two typical subjects will report that the D surface is stiffer than the K surface at various blocks of the experiment. Red, First block; green, second block with no displacement; blue, second block with displacement. The triangles show the estimated probability at each stiffness difference. The lines are the maximum likelihood fits of a sigmoidal function; the horizontal bars on each curve are the 95% confidence interval for the estimation of the point of subjective equality. ***C***, 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 dashed vertical line demonstrates the performance of a "perfect subject" who can accurately estimate whether surface D is stiffer than K. The solid line shows the "regular subject", who would make some mistakes in the transition region (gray shading). A shift in this graph to the left or right, as seen by the dashed lines, would suggest the subject perceives surface D or K, respectively, as stiffer than it really is.

A prediction of the outcome can be made based on Equation 8 and is shown in Figure 4. There, the stiffness profile on the force–position plane during the interaction with an ordinary and shifted surface is depicted. Each solid line in the figure represents a spring surface. The slope of the green line represents the nominal stiffness of a surface, whose boundary is at the origin of the plot. The solid red line has the same slope as the green line, though it is displaced from the origin. If the stiffness were estimated solely based on the maximum force over the actual amount of penetration, there should be no evident change in the perception of stiffness for the displaced surfaces. Assuming that there exists a prior representation of the boundary of the surface (dashed red line), the perceived surface boundary (cyan dot) might be altered by this prior belief. A change in perceived boundary location would affect the perceived penetration and modify the perception of stiffness, as in Equation 8. This manipulation might affect both the reference (K) surface and the displaced (D) surface, as the subject is not aware of which one of the two surfaces is the displaced one. Note that the color of the surface was not uniquely associated with either surface K or D. The association varied randomly from trial to trial. Therefore, on a given trial, the subject was not informed about which surface was actually displaced.

Figure 4*A* shows an underestimation of the stiffness (the slope of the red dashed line is smaller than that of the solid red line) if the boundary is displaced away from the subject; Figure 4*B* shows an overestimated stiffness.

Autoregressive interpretation of the prediction

Autoregressive equations are frequently used in adaptive control algorithms and have been successfully used as a model for motor adaptation during reaching movements (Scheidt et al., 2001; Emken et al., 2007). Here, we use the following autoregressive equation to describe the internal representation of the boundary location:

$$I^m = \omega I^{m-1} + (1-\omega)Y^m, \tag{9}$$

where I is a state variable that describes the perceived penetration and Y is the measured initial position in which the force changes from zero on the mth surface probing (Fig. 1*B*). Y^m is the point at which the force changes from zero on the mth probe. Note that index n is the trial number at which two surfaces are presented to the subject. The m index is the number of the probes into the surfaces. Since subjects probed the surface more than once per trial, the number of probes is larger than the number of trials, as in Figure 1*C*, where a single trial is presented (i.e., n is constant) and 10 probes are made.

Models

We considered three computational models as predictors of the responses made by the subjects.

Nominal penetration. The nominal penetration (NP) model assumes no change in perceived stiffness and therefore implies that the maximum interaction force is divided by the actual penetration into the surface. In this case, Equation 8 changes to

$$K^{\mathrm{NP}} = \frac{\max(\mathrm{Force})}{B}. \tag{10}$$

Constant prior. In the constant prior (CP) model, it is assumed that subjects refer their penetration to a location that remains constant throughout the experiment, i.e., a prior representation is created in the reference block of the experiment (the first block) and from then on, penetration is referred to that constant location. Equation 8 changes to:

$$K^{\mathrm{CP}} = \frac{\max(\mathrm{Force})}{B_{\mathrm{perceived}}} = \frac{\max(\mathrm{Force})}{B + Y_{\mathrm{prior}}}. \tag{11}$$

Variable prior. In the variable prior (VP) model, the penetration length is modified by I, in addition to X_0, which represents the perceived location of the boundary by the subject (Eq. 9). Since the value of I is derived through an autoregressive process, it might change throughout the experiment.

$$K^{\mathrm{VP}} = \frac{\max(\mathrm{Force})}{B_{\mathrm{perceived}}} = \frac{\max(\mathrm{Force})}{B + H - I}. \tag{12}$$

The values of ω in Equation 9 and of Y_{prior} in Equation 11 are fitted to the data as described below. Given the NP, CP, and VP models, a specific value of stiffness for the two surfaces is estimated by each of them. For instance, the estimated stiffness for the D and K surfaces according to the VP model is $K_{\mathrm{D}}^{\mathrm{VP}}$ and $K_{\mathrm{K}}^{\mathrm{VP}}$, respectively.

Derivation of the models

During any given trial, the subjects were presented with two surfaces, which they repeatedly probed. We assumed the subjects assessed the stiffness of each one of the D and K surfaces as $\hat{K}_{\mathrm{D}}$ and $\hat{K}_{\mathrm{K}}$, respectively (Eq. 2). Subjects were then asked to report in a binary fashion which of the two they thought to be stiffer, i.e., to report whether the perceived difference $\hat{\Delta} = \hat{K}_{\mathrm{D}} - \hat{K}_{\mathrm{K}}$ was greater or smaller than zero (Eq. 3). The subjects would report "D" if the stiffness of the D surface was perceived as higher, i.e., $\hat{\Delta} > 0$. To account for noise or uncertainty in the decision process, we assumed that the probability of the subject reporting the D surface as stiffer (a binary variable) had a sigmoidal shape rather than a step function, i.e.,

$$P(\text{subject reporting that } \hat{K}_{\mathrm{D}} > \hat{K}_{\mathrm{K}} | \Delta, \bar{\beta}^{S}) = \frac{1}{1 + e^{-g(\Delta, \bar{\beta}^{S})}}. \tag{13}$$

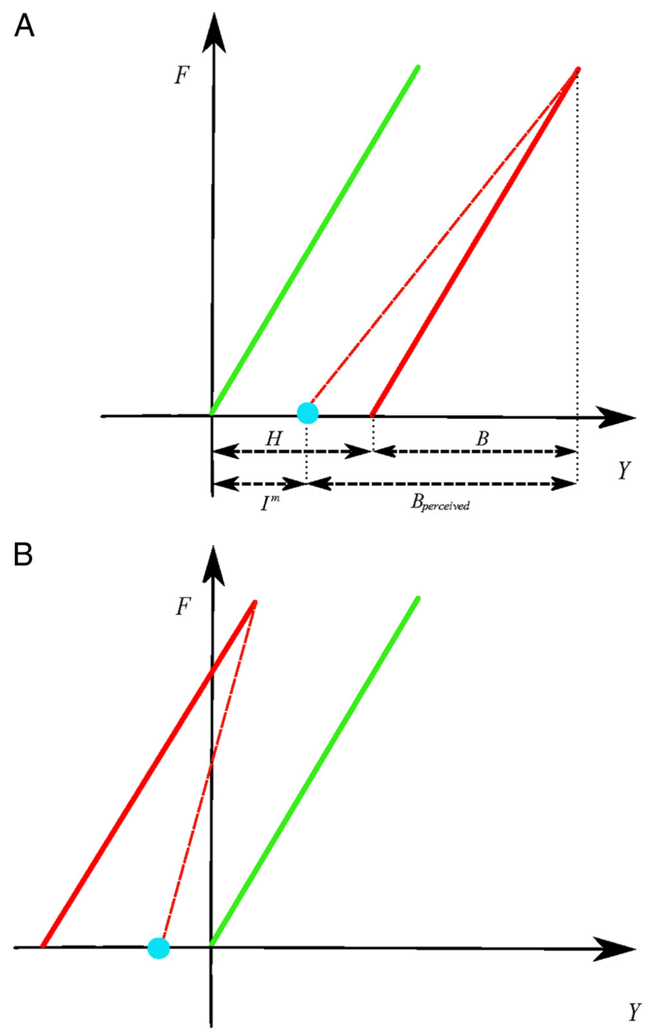

Figure 4. ***A***, ***B***, Prediction of the stiffness change due to a displacement of the surface. ***A***, Case in which the surface was displaced away from the subject. The cyan point represents the perceived position of the boundary; therefore, there is a change in perceived penetration. The solid red and green lines show the nominal stiffness of the surface, which have an equal slope. The slope of the dashed red line represents the perceived stiffness caused by the change in perceived penetration. The change in perceived stiffness is evident, as the slope of the line has become smaller. *B* is the amount of actual penetration and *H* is the displacement in the boundary. Below are the perceived penetration ($B_{\text{perceived}}$) and the assumed location of the boundary (I^m) by the VP model on the *m*th probing movement. ***B***, Same as ***A***, but for the case in which the surface is displaced toward the subject and the slope of the dotted red line becomes greater.

We modeled this function using the logit transformation (Eqs. 5, 6) with the vector of parameters as $\bar{\beta}^S$, which quantifies the uncertainty in subjects' ability to assess stiffness. Furthermore, we assumed the expected value of subject's reports to be Bernoulli distributed with parameter $P(\Delta, \bar{\beta}^S)$ (Ludwig et al., 2005). If we label the binary response of the subject on the *n*th trial as $r_n \in \{0,1\}$, we have $\text{Prob}(r_n = 1|\Delta_n) = P(\Delta_n, \bar{\beta}^S)$.

The expected value for the subject to report that the D surface was stiffer than the K surface (report $r_n = 1$) is $P(\Delta_n, \bar{\beta}^S)$, as in Equation 13. Note that P in Equation 13 is a function of Δ, the actual difference in stiffness between the two surfaces.

Both experiments had three types of trials (see Setup, above): base line, reference, and displaced. Within the baseline trials, the displacement of both the K and D were set to zero, which allowed the assessment of the subject's baseline behavior. Therefore, the vector $\bar{\beta}^S$ was estimated by performing a logistic regression over the data of each subject that was retrieved during the first part of the experiment.

Each of the models (Eqs. 10–12) was used to estimate the apparent stiffness of the two surfaces. For example, with the VP model, we estimated $\tilde{K}_D^{VP}$ and $\tilde{K}_K^{VP}$. The tilde indicates an estimate and the subscript K and D stand for the surfaces estimated on the current trial. For each model, we assumed a level of uncertainty that was subject-specific. That is, we assumed that the expected value of the model reporting the D surface was stiffer than the K surface as follows:

$$P_M(\text{model reports } \tilde{K}_D^{VP} > \tilde{K}_K^{VP} | \tilde{\Delta}, \bar{\beta}^S) = \frac{1}{1 + e^{-g(\tilde{\Delta}, \bar{\beta}^S)}}. \quad (14)$$

In this case, $P(\cdot)$ is a function of $\tilde{\Delta}$, the assumed perceived stiffness difference of $\Delta = K_D^{VP} - K_K^{VP}$.
The uncertainty built in to each model is specific to each subject, due to the dependence on $\bar{\beta}^S$. The equations essentially state that once the difference in stiffness between the two surfaces estimated by the model is fairly close to the difference perceived by the subject, the model is prone to make more errors. This is captured by the fact that for $\tilde{\Delta} = -\beta_1^S/\beta_2^S$ $[g(\tilde{\Delta}) = 0]$, there is a chance level (probability of 0.5) that the model will report the D surface to be stiffer that the K surface.

Once the uncertainty is built into the model, the responses of a model as a function of the actual difference in nominal stiffness (i.e., Δ in Eq. 3) are used to assess the quality or likelihood of the model. We can therefore evaluate Equation 13 for each model as follows:

$$P_M(\text{model reports } \tilde{K}_D^{VP} > \tilde{K}_K^{VP} | \Delta) = \frac{1}{1 + e^{-g(\Delta, \bar{\beta}^M)}}. \quad (15)$$

A parameter vector, $\bar{\beta}^M$, was estimated from the input–output pairs of the above equation (P and Δ) using nonlinear least-square minimization (Bates and Watts, 1988). This process was iterated for each of the models: variable prior ($\bar{\beta}^{MVP}$), constant prior ($\bar{\beta}^{MCP}$), and nominal penetration ($\bar{\beta}^{MNP}$).

The nominal penetration model includes no parameters. The constant prior model includes a single parameter (Y_{prior}), which was estimated from the data of each subject. The estimation of this parameter followed the same lines as the estimation of ω for the VP model (below). That is, we chose the value of Y_{prior} that maximized the likelihood function. For the autoregressive VP model, we sought to establish the likelihood of the responses by the subjects, given the vector $\bar{\beta}^{MVP}$; in this case, the ω parameter was adjusted as in Equation 9. The values that were finally used were those that maximized the likelihood function (minimum of the negative log likelihood value).

The difference between the log likelihood values of two models provides us with a measure of superiority of one model over the other (Kass and Raftery, 1995; Song and Lee, 2006). A difference of 10 is regarded as highly significant. We followed this approach to compare the likelihood value for the three models above and to determine which one better described the data.

Results

Shift in perception

We found that the perception of stiffness was consistently and significantly altered when the object location was unexpectedly shifted, demonstrating the predicted illusion. Occasional displacements caused a change in the perception of stiffness for all 13 subjects and the perceived change was in agreement with the direction of the displacement applied to the object's boundary. In a second experiment, the surface was not fixed to a specific location, but slowly displaced back and forth with a sinusoidal time course (Fig. 2). This slow displacement did not affect the change in perception due to occasional displacement of one of the two surfaces.

Figure 3 shows the estimated probability of reporting that the D surface was stiffer than the K surface as a function of the actual stiffness difference for two typical subjects performing the first experiment. The two central curves correspond to the probability estimate with no displacement in the first (red) and second (green) experimental blocks. The fact that these two curves over-

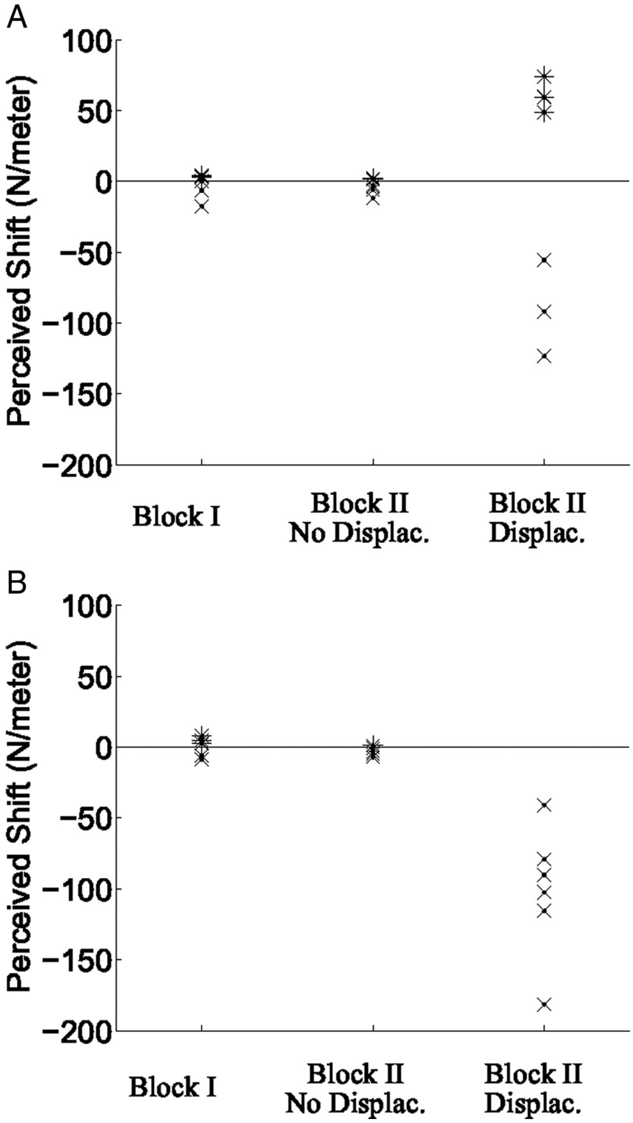

Figure 5. ***A***, ***B***, The amount of shift in the sigmoids for each subject across blocks in newtons per meter for the two experiments (***A***, constant reference experiment; ***B***, moving reference experiment). The Xs correspond to a displacement (Displac.) toward the subjects and asterisks show data for displacements away from the subject.

lap demonstrates the stability of the subjects' responses in this static condition. The blue curves represent the probability estimates in the second experimental block, with displaced boundaries. Figure 3*A* shows the effect on probability estimates of displacing the boundary away from the subject. In this case, the probability curves shifted to the right, indicating that the subjects assessed the displaced surface as being less stiff than the reference K surface. Figure 3*B* shows the effect of displacing the boundary of the D surface toward the subject. In this case, the effect on stiffness estimation was in the opposite direction. The curve shifts to the left, corresponding to an overestimation of the stiffness of the D surface (Fig. 3*C*).

The amount of shift of each logistic curve estimated for each subject is shown in Figure 5. The Xs indicate the data for subjects who experienced a shift toward them; the asterisks correspond to the data of subjects who experienced a displacement away from them. To perform a statistical analysis regarding the significance of the PSE's absolute shift, the values of the PSE for subjects who experienced a backward shift (X) were taken with a negative sign. This would cause all PSE point to have a positive sign (X by flipping the sign; asterisk by construction). Figure 5*A* shows the shifts in the first experiment, in which the baseline conditions were kept constant. Figure 5*B* shows the shifts in the second experiment, in which the boundary slowly moved back and forth with a sinusoidal time–function (for detailed direction of displacement, see Table 1). We tested both scatters for a significant shift in respect to baseline performance. Both had a mean that was statistically different from zero (paired *t* test, $p < 0.05$). This demonstrates that the unexpected displacement of a surface, either stationary or in slow and predictable motion induces, a change in stiffness perception. There was no significant change in the size of displacement between the two experiments (*t* test, $p = 0.19$).

The average amount of exerted force across all subjects in the two experiments was 16.1 ± 4.2 N. The average indentation was 4.27 ± 1.3 cm. There was a significant change in endpoint location during trials in which the boundary was displaced for all subjects (*t* test, $p < 0.001$). An overshoot was evident in cases where the boundary was displaced away from the subject and an undershoot in the other case. Subjects performed as many probing motions as they wished before deciding which surface was more rigid. Our data suggest that the number of probing motions depended upon the similarity between the surfaces, with a significant increase in the number of probes as the stiffness difference decreased (*t* test, $p < 0.021$). The total number of probes was also significantly larger when the stiffness level was more similar (*t* test, $p < 0.028$).

Prediction of the models

The stiffness of an object is defined as the quasistatic interaction force F normalized by the displacement Y inside the object's boundary; that is, $K = F/Y$. Therefore, the perception of stiffness is likely to involve the perception of both tactile information and position information. Studies of manual discrimination of stiffness by Srinivasan and LaMotte (1995) found dissimilarity between the just-noticeable differences for compliant objects with deformable surfaces and rigid surfaces. They concluded that when a surface is deformable, the spatial pressure distribution within the contact region between the hand and the surface is dependent on object compliance. Thus, information from cutaneous mechanoreceptors is sufficient for discrimination of subtle differences in compliance. When the surface is rigid, kinesthetic information is necessary for discrimination, and the discriminability is much poorer than that for objects with deformable boundaries. The perception of stiffness has been addressed by a number of psychophysical studies. Jones and Hunter (1990) had subjects perform a contralateral limb-matching procedure, in which subjects adjusted the stiffness of a motor connected to one arm until it was perceived to be the same as that connected to the other arm. This allowed the authors to assess the fidelity of stiffness perception. Their study shows that the sensitivity to stiffness discrimination is smaller than the sensitivities of force discrimination and displacement discrimination. Understanding how the brain processes haptics will contribute to the development of technological applications, including virtual haptic environments as well as telerobotics, telerehabilitation, and telesurgery, in which the perception of touch must be accurately reconstructed based on the bidirectional transmission of motion and force information (Anderson and Spong, 1989; Hannaford, 1989; Reiner, 2004; Griffiths et al., 2008). As of now, it is not clear which neural mechanisms allow the brain to estimate the rigidity of a contact from different sources of sensory-motor information.

In previous studies (Pressman et al., 2007; Nisky et al., 2008), we introduced a paradigm in which the interaction force with a surface was either lagging or leading the position of the hand. Our goal was to investigate how the relative timing between hand position and force information affects the haptic perception of their ratio, i.e., stiffness. We considered several computational models of stiffness estimation. Some were refuted, but a few were found to generate predictions consistent with reports made by the subjects. All models estimated stiffness by regression of force over hand position inside the object boundary. The relative time shift between these data affects the predictably of the resulting stiffness estimate. We compared different models of this time shift and found that the reports of the subjects were best accounted for by a regression model that matched the maximum interaction force experienced along the trajectory, with the maximum amount of perceived penetration into the surface. Other proxies to the stiffness estimation could be available. However, in this study, we chose to use this model, as it showed the highest agreement with subjects' reports. The rationale for this model is that the effective displacement of the boundary caused by the delay in force reflection changes the perceived amount of penetration into the surface and accordingly changes the perception of stiffness. However, this hypothesis had not yet been put to the test; thus, it is the objective of the present study.

Here, we used peak force for stiffness estimation and considered three alternatives to calculate the perceived penetration.

Nominal penetration

Perceived penetration is simply defined as the actual distance from the boundary of the surface to the furthest point traveled into it. Such a model cannot account for a change in perceived stiffness because of the displacement of the boundary. The NP model has no parameters and assumes that the subjects have perfect information on the each trial regarding the location of the boundary, ignoring information from previous trials.

Constant prior

This model derives the perceived penetration by assuming a fixed location of the boundary throughout the experiment. The model is capable of explaining the change in stiffness perception caused by unexpected shifts of the object's location. The CP model has a single parameter: the subject's estimate of the location of the boundary. This location is assumed to be determined by the subject during the initial familiarization phase. It remains constant through subsequent phases of the experiment.

Variable prior

This model assumes that the object may be moving in the environment in a slow and predictable way. The location of the object boundary is estimated by an autoregressive process that takes into account previous and current locations. The model can account for a constant boundary position (constant reference experiment) and for a slowly moving boundary (moving reference experiment). We used a first order autoregressive (AR) model, which included only one parameter and could be compared with the CP model without risk of overfitting the data.

We compared the predictions of these three models to experimentally derived curves in Figure 6. These predictions were generated so as to mimic the responses of the subject. On a given trial, the model assesses the rigidity of the two surfaces and produces an expected value for the response. By repeating this procedure for various stiffness differences, we obtained the psychometric curves of Figure 6 (for more details, see Materials and Methods, above). The figure compares the prediction of each model with

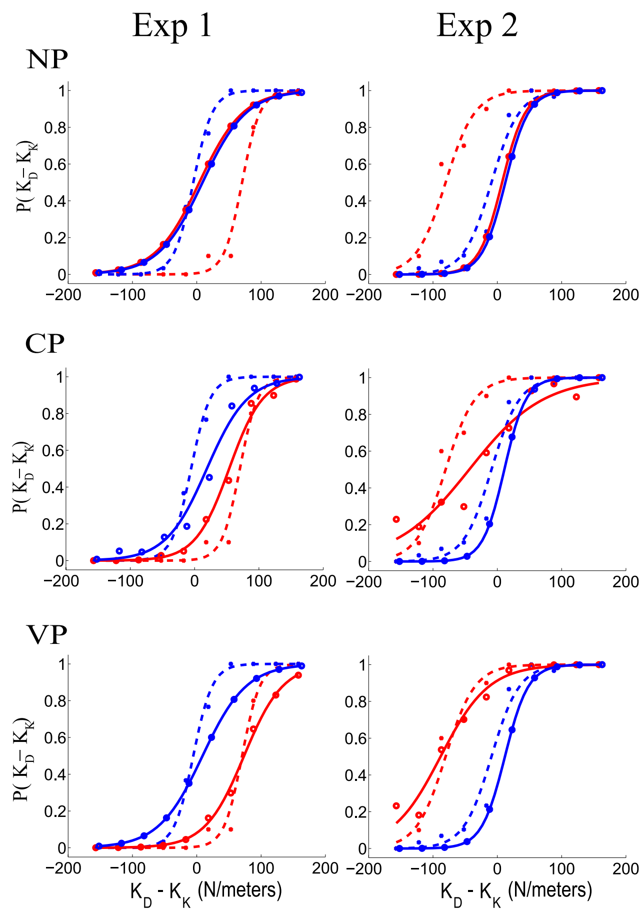

Figure 6. Psychometric curves of a single subject (dashed lines) and the three models (solid lines) for the two experiments (exp). Blue shows data for the nondisplaced trials and red shows displaced trials. The red lines were slightly shifted to the left along the *x*-axis to aid visualization of both curves.

the performance of a single subject from each group (two different subjects). The curves show the estimated frequency with which a subject/model indicates that the D surface was stiffer than the K surface as a function of the actual stiffness difference ($\Delta = K_D - K_K$). The dashed lines (red and blue) represent the subject's responses and the solid lines are the predictions of the models. The red lines (dashed and solid) show the data for the trials in which a displacement took place; the blue lines are for trials with no displacement.

It is evident that the NP model cannot capture the behavior of the subject in either experiment. The CP and VP models both show some degree of agreement. Although the VP model appears to better account for responses made by the subjects, the comparison between the two models must be based on more rigorous analysis. We therefore performed a likelihood ratio test. The likelihood ratio estimates how likely one model is to be correct compared with another, given some dataset. In our case, the dataset contained the responses of the subject and the actual stiffness difference in each trial (MacKay, 2002; Körding et al., 2007).

We compared the likelihood of the observed data given each model, $P(r|\Delta, \bar{\beta})$, where r is the subject's response that the displaced surface is stiffer, Δ is the actual difference in stiffness and $\bar{\beta}$ is model's parameters. A judgment about the validity of the CP and VP models is derived by comparing the likelihood of the observed dataset under each model. We consider the model that shows a higher log-likelihood value to be more accurate. This is actually the lowest value in Table 1, which displays negative log likelihood values (MacKay, 2002). More specifically, Table 1 shows the estimate of the two values $L(\bar{\beta}^{\mathrm{MVP}})$ and $L(\bar{\beta}^{\mathrm{MCP}})$ for all subjects in both experiments for both CP and VP models.

The VP model was selected as the simplest AR model with one parameter, $I^p = \omega I^{m-1} + (1 - \omega)Y^m$ (Eq. 9). The results for this model are shown in Table 1. A higher-order VP model (order three) and the order-one model had similar likelihood values across subjects (mean differences of 0.7 and 0.2 for constant reference and variable reference, respectively, which imply no significant effect) (Kass and Raftery, 1995).

Table 1 shows the difference for the two experiments, along with the mean value for each experiment. For 12 of 13 subjects, across both experiments, the VP model yielded a better likelihood score than the CP model. Following the guidelines regarding likelihood ratios (or differences if logistic functions were used) given in Kass and Raftery (1995), we can see that for the first experiment, the VP model is substantially better (2–6) than the constant prior, whereas on the second experiment it is decisive ($>$10).

The predictions of the models regarding the slope of the sigmoid curves were also compared with the subjects' data and were found to significantly underestimate the actual slopes (predicting a smaller slope, *t* test, $p < 0.01$). Since the models were derived from the subject's uncertainty during the first part (Eqs. 14, 15), it follows that the uncertainty should grow in such a case by construction of the model. To keep our model as simple as possible, we did not introduce additional parameters. Therefore, we restrict our prediction to the shift in the curve.

Noise hypothesis

In Materials and Methods, above, and, more specifically, in Equation 13, we introduced the idea of including the uncertainty of subjects' responses as a form of noise in the computational model. By using the uncertainty of subjects' responses within the initial part of the experiment (baseline), we implicitly assumed that there was constant uncertainty throughout the experiment. In a logistic curve, the uncertainty of the responses is captured by the slope of the sigmoid at the 0.5 probability level. This slope is expressed by β_2 parameter in Equation 6. To test the hypothesis that the overall amount of uncertainty did not change, the β_2 values were compared between baseline trials and test trials in which a displacement took place. There was no significant change in the mean value across all subjects (*t* test, $p > 0.5$; mean reference, 0.045 ± 0.018 and mean displaced, 0.049 ± 0.023, across all subjects in both experiments. Each experiment showed the same trend).

Discussion

Psychophysical studies have revealed a multitude of perceptual distortions that are connected to our estimate of primary visual features of objects, such as length, orientation, and colors. Here, we have investigated the perception of mechanical properties as subjects formed them through object manipulation. We asked subjects to touch the boundary of mechanically simulated virtual objects and asked them to report their perception of the object's stiffness.

We found that an occasional displacement of an otherwise stationary (or slowly varying) surface leads to an illusion of change in stiffness of the surface. The displacement can be either towards or away from the subject, resulting in an increase or decrease of perceived stiffness, respectively.

Our findings suggest that an adaptive mechanism based on combining current information and prior knowledge may have a direct impact on the perception of the mechanical properties of the environment we come in contact. This appears to be a direct consequence of the observations that the perception of an object's rigidity arises from a process analogous to a regression of force over position information, and that the information about the location of the object boundary is part and parcel of a mechanism of probabilistic state estimation.

We considered three computational models to account for the responses of the subjects. The first, the nominal penetration model, simply uses the maximum interaction force during the penetration of the hand into the surface divided by the extent of penetration (*F*/*Y*). The next two models included a correction to the extent of penetration, which accounts for the change in perception of length due to prior knowledge. The CP model assumes constant representation of the object location whereas the VP model deploys an autoregressive process to model the memory effect acquired on previous trials. Our results clearly refute the nominal penetration model. The likelihood values for the VP model show higher plausibility for such a model with respect to the constant penetration model. This is especially evident in the second experiment, where the boundary of the two surfaces was slowly varied throughout the task.

In this work, we consider a simple model, which only accounts for the bias in stiffness perception and does not address the issue of noise within this process. A more general model would include a term that describes not only the shift in perception but also the variability in assessment of stiffness. We have verified that adding the variance of a normally distributed noise does not change the results about the bias in perception (i.e., the mean of this normally distributed shift) and its dependence on prior. Therefore, we used the simplest model that can explain what we see as the essence of this illusion—the overall shift in perceived stiffness.

The variable prior model described above is the preferred model in this case. A key issue with such a model is the single reference point or the single prior used throughout the probing activity (Eq. 9). The alternative of two reference points, or two priors (one for each surface), holds in the implicit assumption that the nervous system can detect each of the surfaces and associate them to the appropriate prior. This is contradicted by the very existence of the illusion, since each surface would be associated to its prior and no misperception would be created. Thus, our findings and the Occam's razor argument support a simple model with a single prior for all surfaces.

Theoretically, it would have been appealing to design an experimental condition without prior, in which the illusion is eliminated. In practice, the experimental removal of a prior is quite difficult. If one were to present the objects in different locations, one could still not rule out that a prior corresponding to the average location is not formed. In our experiment, the prior formation seemed to occur rapidly. A prior had been formed after a single probing motion. Moreover, our workspace was limited by the robotic manipulandum; therefore, the illusion was robust and seemed to appear in all the possible conditions of our setup. Future studies should also consider the transfer to the other hand or to completely different configurations of the arm. However, the existence or absence of such transfer does not undermine our hypothesis.

There are similarities between the boundary–stiffness illusion and two other notable haptic illusions: the size–weight illusion (SWI) and the material–weight illusion (MWI). In the SWI (Charpentier, 1891; Chouinard et al., 2009; Brayanov and Smith, 2010), subjects perceive smaller objects as being heavier. In the MWI, objects that appear to be made of lightweight materials such as Styrofoam are perceived to be lighter than objects of equal weight that appear to be made of heavier materials such as aluminum (Ellis and Lederman, 1999; Buckingham et al., 2009). In both illusions, although perception is swayed by appearances, the motor control system learns rapidly to operate correctly. So, for example, as subjects repeat lifting two objects that have the same weight but different sizes, they quickly learn to apply the same grip forces. However, when asked about the objects' weight, they persist in reporting that the smaller object is heavier. Flanagan and Beltzner (2000) considered this finding as evidence "that the cognitive/perceptual system can operate independently of the sensorimotor system." Buckingham et al. (2009) recently reported a similar discrepancy between motor and perceptual responses in the MWI. Brayanov and Smith (2010) have described the perceptual effect in the SWI as an anti-Bayesian phenomenon. They suggest that in the SWI there is an enhancement of the unexpected violation of the prior expectation that the smaller objects weighs less than the large object. Chouinard et al. (2009) suggested that the SWI is connected to the perception of density; the smaller object is considered to have a higher density than the large object. This interpretation is supported by the results of an fMRI study, where the ventral premotor area shows activity while subjects experience the SWI. The same area appears to integrate weight and size information to represent the density of objects. In contrast, no activity was observed during SWI in sensory areas or in M1, which appear to provide representations for size and weight. As for SWI and MWI, the illusion reported here is related to a discrepancy between prior expectation and sensory haptic input. However, in our case, the prior does not arise from visual information but from the proprioceptive experience of the arm configuration at the point of contact with the virtual surface. In our stiffness illusion, rather than a bias induced by the experienced object properties, the perceptual effect appears to be directly connected with sensory-motor behavior. In our case, a shift forward or backward of the boundary results in wrong estimate of hand position at contact and this is reflected in a forward or backward, respectively, shift of the whole probing motion. However, the effect on stiffness perception suggests that this shift is not accounted for correctly when the stiffness is estimated by regression of force over penetration. On one hand, the experience of peak force is almost unaffected by the perturbation of the boundary. On the other hand, however, a forward shift of the boundary leads to an overestimation of the amount of movement of the hand and a backward shift leads to an underestimation. This error in estimation of the state of motion of the hand appears to be the prime cause of the stiffness illusion. Therefore, in the stiffness illusion, we would expect to observe a greater amount of correlation between activities in sensorimotor and perceptual areas compared with the size-weight and material-weight illusions.

These findings support the general view that our ability to assess the mechanical properties of the objects that we come in contact with depend on our nervous system combining state and force information (Venkadesan and Valero-Cuevas, 2008) and carrying out a process that is analogous to statistical regression. Sudden and unexpected displacements of the contact point result in predictable effects on the perceived impedance. Earlier studies (Ernst and Banks, 2002; Körding and Wolpert, 2004) showed that the nervous system combines prior information with current sensory evidence to estimate the state of motion (position and velocity) of the limbs. Our autoregressive model (VP) is consistent with this probabilistic framework; it postulates that subjects were constantly estimating the location of the boundary based on a combination of current evidence and a priori information from the past experience of contacts. Here, we have shown for the first time the consequence of such state estimation on the assessment of mechanical impedance—the stiffness—at the point of interaction with the environment.

References

Anderson RJ, Spong MW (1989) Bilateral control of teleoperators with time delay. IEEE Trans Automat Contr 34:494–501.

Bates DM, Watts DG (1988) Nonlinear regression analysis and its applications: Wiley series in probability and statistics. New York: Wiley.

Botvinick M, Cohen J (1998) Rubber hands 'feel' touch that eyes see. Nature 391:756.

Brayanov JB, Smith MA (2010) Bayesian and "anti-Bayesian" biases in sensory integration for action and perception in the size-weight illusion. J Neurophysiol 103:1518–1531.

Buckingham G, Cant JS, Goodale MA (2009) Living in a material world: how visual cues to material properties affect the way that we lift objects and perceive their weight. J Neurophysiol 102:3111–3118.

Charpentier A (1891) Analyse experimentale quelques elements de la sensation de poids. Arch Physiol Normales Pathologiques 3:122–135.

Chouinard PA, Large ME, Chang EC, Goodale MA (2009) Dissociable neural mechanisms for determining the perceived heaviness of objects and the predicted weight of objects during lifting: an fMRI investigation of the size-weight illusion. Neuroimage 44:200–212.

Conditt MA, Gandolfo F, Mussa-Ivaldi FA (1997) The motor system does not learn the dynamics of the arm by rote memorization of past experience. J Neurophysiol 78:554–560.

Ellis RR, Lederman SJ (1999) The material-weight illusion revisited. Percept Psychophys 61:1564–1576.

Emken JL, Benitez R, Sideris A, Bobrow JE, Reinkensmeyer DJ (2007) Motor adaptation as a greedy optimization of error and effort. J Neurophysiol 97:3997–4006.

Ernst MO, Banks MS (2002) Humans integrate visual and haptic information in a statistically optimal fashion. Nature 415:429–433.

Flanagan JR, Beltzner MA (2000) Independence of perceptual and sensorimotor predictions in the size-weight illusion. Nat Neurosci 3:737–741.

Geldard FA, Sherrick CE (1972) The cutaneous "rabbit": a perceptual illusion. Science 178:178–179.

Gescheider GA (1997) Psychophysics: the fundamentals, 3 Ed. Mahwah, NJ: Lawrence Erlbaum.

Griffiths PG, Gillespie RB, Freudenberg JS (2008) A fundamental tradeoff between performance and sensitivity within haptic rendering. IEEE Trans Rob 24:537–548.

Halligan PW (2002) Phantom limbs: the body in mind. Cogn Neuropsychiatry 7:251–269.

Hannaford B (1989) A design framework for teleoperators with kinesthetic feedback. IEEE Trans Rob Autom 5:426–434.

Hosmer DW, Lemeshow S (2000) Applied logistic regression: Wiley series in probability and statistics. New York: Wiley.

Jones LA (1988) Motor illusions: what do they reveal about proprioception? Psychol Bull 103:72–86.

Jones LA, Hunter IW (1990) A perceptual analysis of stiffness. Exp Brain Res 79:150–156.

Kanizsa G (1955) Margini quasi-percettivi in campi con stimolazione omogenea. Rivista di Psicologia 49:7–30.

Kass RE, Raftery AE (1995) Bayes factors. J Am Stat Assoc 90:773–795.

Körding KP, Wolpert DM (2004) Bayesian integration in sensorimotor learning. Nature 427:244–247.

Körding KP, Beierholm U, Ma WJ, Quartz S, Tenenbaum JB, Shams L (2007) Causal inference in multisensory perception. PLoS One 2:e943.

Levitin DJ (2002) Foundations of cognitive psychology: core readings. Cambridge, MA: MIT.

Ludwig CJ, Gilchrist ID, McSorley E, Baddeley RJ (2005) The temporal impulse response underlying saccadic decisions. J Neurosci 25:9907–9912.

MacKay DJ (2002) Information theory, inference and learning algorithms. Cambridge: Cambridge UP.

McGurk H, MacDonald J (1976) Hearing lips and seeing voices. Nature 264:746–748.

Munhall KG, Gribble P, Sacco L, Ward M (1996) Temporal constraints on the McGurk effect. Percept Psychophys 58:351–362.

Nisky I, Mussa-Ivaldi FA, Karniel A (2008) A regression and boundary-

crossing-based model for the perception of delayed stiffness. IEEE Trans Haptics 1:73–82.
Ohnishi H, Mochizuki K (2007) Effect of delay of feedback force on perception of elastic force: a psychophysical approach. IEICE Trans Comm 90-B:12–20.
Peled A, Pressman A, Geva AB, Modai I (2003) Somatosensory evoked potentials during a rubber-hand illusion in schizophrenia. Schizophr Res 64:157–163.
Pressman A, Welty LJ, Karniel A, Mussa-Ivaldi FA (2007) Perception of delayed stiffness. Int J Robotic Res 26:1191–1203.
Pressman A, Nisky I, Karniel A, Mussa-Ivaldi FA (2008) Probing virtual boundaries and the perception of delayed stiffness. Adv Robotic 22:119–140.
Reiner M (2004) The role of haptics in immersive telecommunication environments. IEEE Trans Circ Syst Vid Technol 14:392–401.
Renier L, Volder AGD (2005) Cognitive and brain mechanisms in sensory substitution of vision: a contribution to the study of human perception. J Integr Neurosci 4:489–503.
Scheidt RA, Dingwell JB, Mussa-Ivaldi FA (2001) Learning to move amid uncertainty. J Neurophysiol 86:971–985.
Shadmehr R, Mussa-Ivaldi FA (1994) Adaptive representation of dynamics during learning of a motor task. J Neurosci 14:3208–3224.
Solso RL (2001) Cognitive psychology. Boston: Allyn and Bacon.
Song XY, Lee SY (2006) Model comparison of generalized linear mixed models. Stat Med 25:1685–1698.
Srinivasan MA, LaMotte RH (1995) Tactual discrimination of softness. J Neurophysiol 73:88–101.
Venkadesan M, Valero-Cuevas FJ (2008) Neural control of motion-to-force transitions with the fingertip. J Neurosci 28:1366–1373.
Wichmann FA, Hill NJ (2001a) The psychometric function. I. Fitting, sampling, and goodness of fit. Percept Psychophys 63:1293–1313.
Wichmann FA, Hill NJ (2001b) The psychometric function. II. Bootstrap-based confidence intervals and sampling. Percept Psychophys 63:1314–1329.