IEEE Robotics and Automation Letters (RAL) paper presented at the
2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)
October 25-29, 2020, Las Vegas, NV, USA (Virtual)

Hybrid Force-Moment Braking Pulse: A Haptic Illusion to Increase the Perceived Hardness of Virtual Surfaces

Ashkan Pourkand[1] and Jake J. Abbott[2]

***Abstract*— A perennial challenge when rendering a virtual surface with an impedance-type haptic interface is making the surface feel hard without destroying its realism, since simply increasing its stiffness can lead to instability. One way to increase the perceived hardness without increasing stiffness is to implement a braking pulse or other high-frequency haptic contact event. Traditionally, such events are implemented as a force along the surface normal, which may leave some of the actuators of the haptic device underutilized. We propose a hybrid force-moment braking pulse, which includes a nonrealistic rendered moment to exploit a haptic illusion. We describe how to implement such a hybrid force-moment braking pulse in general, considering the saturation of the haptic device's actuators. In a human-subject study, we find that a virtual surface rendered with these hybrid force-moment braking pulses is perceived as harder than the same virtual surface rendered with a traditional braking pulse, without harming the surface's realism, for the majority of users. The moment-based haptic illusion also has the potential to be superimposed on other types of haptic contact events to improve the perceived hardness.**

I. Introduction

With impedance-type haptic interfaces—characterized by low inertia, low friction, and backdrivability—rendering of virtual surfaces is typically based on the implementation of a unilateral stiffness that penalizes penetration into the surface, with a force $\boldsymbol{f} = \hat{\boldsymbol{n}}$kx (units N) proportional to the penetration depth x (units m) via a stiffness k (units N/m), where $\hat{\boldsymbol{n}}$ is the surface normal at the point of contact (i.e., x is measured in the $-\hat{\boldsymbol{n}}$ direction). A perennial challenge when implementing such virtual surfaces is making them feel simultaneously realistic and hard. Here, "hardness" refers to a subjective perception of a surface, as opposed to alternate definitions used within the materials and solid-mechanics communities. Increasing the stiffness of a virtual surface is the most straightforward way to make it feel harder. However, it is well known that increasing the stiffness too high leads to nonpassive and unstable behaviors (e.g., vibrations) that destroy the illusion of the intended passive surface [1]–[3].

The "rate hardness" of a contact event with a virtual surface, which is the ratio of the initial rate of change in force to the initial penetration velocity, with units (N/s)/(m/s) = N/m, is known be a good measure of how hard the surface feels [4]. There is evidence that "extended rate hardness", which replaces the *initial* rate of change in force with the *maximum* rate of change in force, has even better correlation with perceived hardness [5].

One way to increase the (extended) rate hardness of a virtual surface without increasing its stiffness is to implement a braking pulse [6]. A braking pulse is a subclass of event-based haptic feedback cues designed to improve the transient sensations perceived during a tool-mediated contact event with a virtual surface without negatively affecting the stability of the surface [7], [8]. In their most basic form, braking pulses are implemented as force impulses designed to cancel the momentum of the hand and stylus when they impact the virtual surface. Let $\boldsymbol{p} = m\boldsymbol{v}$ be the incoming momentum (units N·s), where m (units kg) is the estimated mass of the hand and stylus and $\boldsymbol{v}$ is the velocity (units m/s) of the haptic interaction point (HIP) during the contact event. The braking pulse is typically implemented as an impulse spread over N cycles of the haptic update period T (typically $T = 0.001$ s) with

$$\boldsymbol{f} = -\left(\frac{1}{NT}\right)\hat{\boldsymbol{n}}\hat{\boldsymbol{n}}^{\top}\boldsymbol{p} \tag{1}$$

with N chosen as small as possible so that the braking-pulse event is perceived as instantaneous. The maximum magnitude of the force that can be rendered will be limited in any real system due to saturation of the amplifier or power supply driving the motors. This will fundamentally limit the achievable hardness.

However, during such a saturation event, not every actuator saturates simultaneously, and typically only a single actuator saturates. This leaves the remaining actuators potentially underutilized. Consider the case of a contact event with a six-degree-of-freedom (6-DOF) Phantom Premium. The simple kinematic structure of this haptic device makes it easy to visualize the phenomenon of interest here. To render the desired upward force at the HIP in the configuration shown in Fig. 1(a), only a single motor (Motor 2) is used. We see that the user experiences both a reaction force and a reaction moment to balance the rendered force quasistatically. That is, both forces and moments at the hand are part of the typical haptic experience of tool-mediated contacts. If we were to superimpose a torque at Motor 5, as shown in Fig. 1(b), we could increase the magnitude of the reaction moment at the hand without affecting its direction, and without affecting the reaction force. It is our hypothesis that, if rendered over a short duration, this will lead to a perception of increased hardness without harming the realism of the rendering. In this

*This work was supported by the National Science Foundation under grant number 1423273.

[1]A. Pourkand is with the School of Computing and the Robotics Center, University of Utah, Salt Lake City, UT 84112 USA ashkan.pourkand@utah.edu

[2]J. J. Abbott is with the Department of Mechanical Engineering and the Robotics Center, University of Utah, Salt Lake City, UT 84112 USA jake.abbott@utah.edu

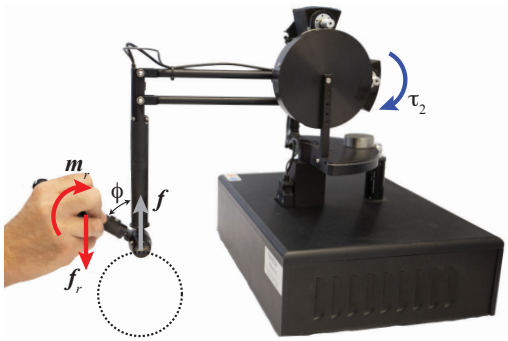

(a) Traditional force braking pulse.

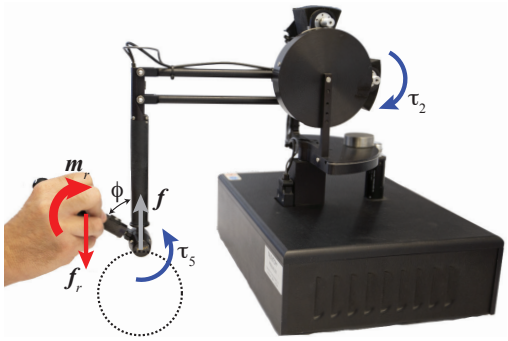

(b) Hybrid force-moment braking pulse.

Fig. 1: Comparison of braking-pulse methods, depicted with a 6-DOF Phantom Premium. The gray arrow is the force to be rendered at the HIP, caused by an impact with a virtual sphere. Blue arrows indicate actuator torques. Red arrows indicate the reaction force and moment of the hand (assuming static equilibrium). Adding τ_5 increases the magnitude of $\boldsymbol{m}_r$.

paper, we describe how to implement such a hybrid force-moment braking pulse in general, considering the saturation of the haptic device's actuators, and we demonstrate the existence of this haptic illusion [9] through human-subjects studies.

Our idea is not without precedent. Prior works have used a gyroscope or a moving center of mass in ungrounded kinesthetic haptic displays to generate kinesthetic moment illusions [10]–[12]. Other works have shown, in the context of high-frequency vibrotactile display on an ungrounded stylus, how moment channels can be used to increase the perceived sensation magnitude, without users being able to distinguish moments from forces [13]–[15]. It is also known that cutaneous feedback has a significant effect on the perception of hardness [16]. Because a braking pulse is fundamentally a "high-frequency" event, albeit of a nonperiodic nature, our hypothesis in this study is that underutilized moment actuation can be used to increase the perceived magnitude of braking pulses, which will ultimately lead to virtual surfaces that appear harder than with traditional braking pulses, but without any explicit attempt to increase the rate hardness.

The case of the 6-DOF Phantom Premium depicted in Fig. 1 also enables us to visualize the limitations of the proposed method. The reaction moment experienced at the hand is fundamentally due to the moment arm from the HIP to hand. Let ϕ represent the positive angle between the surface normal and the axis of the stylus. As ϕ is reduced, so is the reaction moment at the hand. For haptic interactions with $\phi = 0$ (e.g., needle puncture), we would not expect any reaction moment at the hand, and thus would not expect a torque applied at Motor 5 to contribute to the illusion of a larger puncture force.

II. Force-Moment Illusion in Tool-Mediated Contact Events

When a static force $\boldsymbol{f}$ is rendered at the HIP, simulating a typical tool-mediated contact with an environment, the user experiences a reaction force

$$\boldsymbol{f}_r = -\boldsymbol{f} \tag{2}$$

and reaction moment

$$\boldsymbol{m}_r = \boldsymbol{r} \times \boldsymbol{f} = \mathbb{S}\{\boldsymbol{r}\}\boldsymbol{f} \tag{3}$$

as depicted in Fig. 1(a), where $\boldsymbol{r}$ is the vector from the HIP to the grasp, and $\mathbb{S}\{\boldsymbol{r}\}$ is the skew-symmetric matrix packing of $\boldsymbol{r}$ that represents the cross-product operation.

When a static moment $\boldsymbol{m}$ is rendered at the HIP, the user experiences a reaction moment

$$\boldsymbol{m}_r = -\boldsymbol{m}, \tag{4}$$

which has no resulting reaction force. When a static force $\boldsymbol{f}$ and moment $\boldsymbol{m}$ are simultaneously rendered at the HIP as depicted in Fig. 1(b), the user experiences a reaction force $\boldsymbol{f}_r$ and reaction moment $\boldsymbol{m}_r$ that are a combination of (2)–(4):

$$\begin{bmatrix} \boldsymbol{f}_r \\ \boldsymbol{m}_r \end{bmatrix} = \begin{bmatrix} -\mathbb{I} & \mathbb{O} \\ \mathbb{S}\{\boldsymbol{r}\} & -\mathbb{I} \end{bmatrix} \begin{bmatrix} \boldsymbol{f} \\ \boldsymbol{m} \end{bmatrix} \tag{5}$$

where $\mathbb{I}$ is the identity matrix and $\mathbb{O}$ is a zero matrix. This situation is unlike the real tool-mediated interactions that humans are accustomed to, since there is typically no applied moment during a contact event. It is our hope that this can be exploited in a haptic illusion [9]. It is our hypothesis that rendering a moment

$$\boldsymbol{m} = \mathbb{S}\{\boldsymbol{r}\}(\boldsymbol{f}_{\text{des}} - \boldsymbol{f}) \tag{6}$$

during a contact event (i.e., if this moment is applied over a very short duration) can partially create the illusion of a desired force $\boldsymbol{f}_{\text{des}}$ in cases in which it is not possible to render $\boldsymbol{f} = \boldsymbol{f}_{\text{des}}$ directly. Here, $\boldsymbol{f}_{\text{des}} - \boldsymbol{f}$ can be viewed as the shortfall in force capability. Such a method would likely be conservative (i.e., the perceived force magnitude would likely still be less than $\|\boldsymbol{f}_{\text{des}}\|$) since the resulting reaction force will be less than what would be expected with $\boldsymbol{f} = \boldsymbol{f}_{\text{des}}$ and $\boldsymbol{m} = \boldsymbol{0}$.

It should be noted that this haptic illusion does not work quasistatically. An applied moment at the HIP is perceived as physically unrealistic of a simply contact with a surface.

III. Kinematics of the Haptic Device

As we survey 6-DOF impedance-type haptic devices, we find that they are typically constructed as open-chain (i.e., serial) mechanisms with six actuators (e.g., Phantom Premium) or as closed-chain (i.e., parallel) mechanisms with more than six actuators (e.g., Entact W6D, Quanser HD^2). In robots with an open-chain design, it is common to quantify the robot's configuration-dependent manipulability using a Jacobian of the robot's forward kinematics:

$$\begin{bmatrix} \boldsymbol{v} \\ \boldsymbol{\omega} \end{bmatrix} = \mathbb{J}_{\text{fk}}\{\boldsymbol{\theta}\}\dot{\boldsymbol{\theta}}, \tag{7}$$

where $\boldsymbol{v}$ and $\boldsymbol{\omega}$ are the velocity and angular-velocity vectors of the end-effector, respectively, and $\boldsymbol{\theta} = [\theta_1 \ \cdots \ \theta_n]^\top$ is the $n \times 1$ array of joint angles [17]. However, we find that using a Jacobian of the haptic device's inverse kinematics [18],

$$\dot{\boldsymbol{\theta}} = \mathbb{J}_{\text{ik}}\{\boldsymbol{\theta}\}\begin{bmatrix} \boldsymbol{v} \\ \boldsymbol{\omega} \end{bmatrix}, \tag{8}$$

generalizes well for most haptic devices of interest, including both the 6-DOF Phantom Premium discussed previously and the Entact W6D used in the subsequent experiment. Using the principal of virtual work, the force $\boldsymbol{f}$ and moment $\boldsymbol{m}$ that must be applied at the HIP to balance a set of joint torques $\boldsymbol{\tau} = [\tau_1 \ \cdots \ \tau_n]^\top$, in static equilibrium, are calculated as

$$\begin{bmatrix} \boldsymbol{f} \\ \boldsymbol{m} \end{bmatrix} = \mathbb{J}_{\text{ik}}\{\boldsymbol{\theta}\}^\top \boldsymbol{\tau}, \tag{9}$$

and the joint torques to achieve a desired force and moment at the HIP can be calculated using the pseudoinverse:

$$\boldsymbol{\tau} = \left(\mathbb{J}_{\text{ik}}\{\boldsymbol{\theta}\}^\top\right)^\dagger \begin{bmatrix} \boldsymbol{f}_{\text{des}} \\ \boldsymbol{m}_{\text{des}} \end{bmatrix}. \tag{10}$$

Each of the actuators has a maximum torque magnitude, collected in the array $\boldsymbol{\tau}_{\max} = [\tau_{\max,1} \ \cdots \ \tau_{\max,n}]^\top$. A solution to (10) is valid (i.e., the desired force and moment are achievable) if the magnitude of each element in $\boldsymbol{\tau}$ is not greater than the respective element in $\boldsymbol{\tau}_{\max}$.

IV. Hybrid Force-Moment Braking Pulse

Let us assume that we have a desired (force-only) braking pulse $\boldsymbol{f}_{\text{des}}$, determined using (1). We begin by computing the joint torques, $\boldsymbol{\tau}_{\boldsymbol{f}}$, required to render $\boldsymbol{f}_{\text{des}}$, using (10) with $\boldsymbol{m}_{\text{des}} = \boldsymbol{0}$. If $\boldsymbol{\tau}_{\boldsymbol{f}}$ is achievable (i.e., if each element in $\boldsymbol{\tau}_{\boldsymbol{f}}$ has a magnitude that does not exceed its respective maximum value stored in $\boldsymbol{\tau}_{\max}$), then we have found a solution to generate the desired braking pulse, and utilizing the force-moment illusion is unnecessary; this is the first possible outcome.

If $\boldsymbol{\tau}_{\boldsymbol{f}}$ is not achievable, then we can scale the requested force linearly by a constant $0 < \gamma < 1$ to ensure that no joint torque is above its maximum limit, such that at least one joint torque in $\gamma\boldsymbol{\tau}_{\boldsymbol{f}}$ will be at its maximum limit. If this is the case, our rendered braking pulse $\gamma\boldsymbol{f}_{\text{des}}$ will have a shortfall from what was desired from (1). We would like to make up for this shortfall with an equivalent moment, if possible, using (6), which takes the form

$$\boldsymbol{m}_{\text{des}} = \mathbb{S}\{\boldsymbol{r}\}(\boldsymbol{f}_{\text{des}} - \gamma\boldsymbol{f}_{\text{des}}) = (1-\gamma)\mathbb{S}\{\boldsymbol{r}\}\boldsymbol{f}_{\text{des}} \tag{11}$$

Algorithm 1 Hybrid Force-Moment Braking Pulse

procedure CALCULATE JOINT TORQUES($\boldsymbol{f}_{\text{des}}$,$\boldsymbol{r}$)
  $\boldsymbol{\tau}_{\boldsymbol{f}} \leftarrow (\mathbb{J}_{\text{ik}}^\top)^\dagger[\boldsymbol{f}_{\text{des}}^\top \ \boldsymbol{0}^\top]^\top$
  $\boldsymbol{\tau}_{\text{maxHigh}} \leftarrow \boldsymbol{\tau}_{\max}$
  $\boldsymbol{\tau}_{\text{maxLow}} \leftarrow -\boldsymbol{\tau}_{\max}$
  $\gamma \leftarrow$ SCALE $(\boldsymbol{\tau}_{\text{maxHigh}}, \boldsymbol{\tau}_{\text{maxLow}}, \boldsymbol{\tau}_{\boldsymbol{f}})$
  if γ = 1 **then**
    return $\boldsymbol{\tau} \leftarrow \boldsymbol{\tau}_{\boldsymbol{f}}$ ▷ Outcome 1
  else
    $\boldsymbol{m}_{\text{des}} \leftarrow (1-\gamma)\mathbb{S}\{\boldsymbol{r}\}\boldsymbol{f}_{\text{des}}$
    $\boldsymbol{\tau}_{\boldsymbol{m}} \leftarrow (\mathbb{J}_{\text{ik}}^\top)^\dagger[\boldsymbol{0}^\top \ \boldsymbol{m}_{\text{des}}^\top]^\top$
    $\boldsymbol{\tau}_{\text{maxHigh}} \leftarrow \boldsymbol{\tau}_{\text{maxHigh}} - \gamma\boldsymbol{\tau}_f$
    $\boldsymbol{\tau}_{\text{maxLow}} \leftarrow \boldsymbol{\tau}_{\text{maxLow}} - \gamma\boldsymbol{\tau}_f$
    $\delta \leftarrow$ SCALE $(\boldsymbol{\tau}_{\text{maxHigh}}, \boldsymbol{\tau}_{\text{maxLow}}, \boldsymbol{\tau}_{\boldsymbol{m}})$
    return $\boldsymbol{\tau} \leftarrow \gamma\boldsymbol{\tau}_{\boldsymbol{f}} + \delta\boldsymbol{\tau}_{\boldsymbol{m}}$ ▷ Outcome 2

procedure SCALE($\boldsymbol{x}_{\text{maxHigh}}$,$\boldsymbol{x}_{\text{maxLow}}$,$\boldsymbol{x}$)
  $s \leftarrow 1$
  for *i:*=1 **to** length($\boldsymbol{x}$) **do**
    if $\boldsymbol{x}(i) > 0$ **and** $\boldsymbol{x}_{\text{maxHigh}}(i)/\boldsymbol{x}(i) < \gamma$ **then**
      $s \leftarrow \boldsymbol{x}_{\text{maxHigh}}(i)/\boldsymbol{x}(i)$
    if $\boldsymbol{x}(i) < 0$ **and** $\boldsymbol{x}_{\text{maxLow}}(i)/\boldsymbol{x}(i) < \gamma$ **then**
      $s \leftarrow \boldsymbol{x}_{\text{maxLow}}(i)/\boldsymbol{x}(i)$
  return s

We compute the joint torques, $\boldsymbol{\tau}_{\boldsymbol{m}}$, required to render $\boldsymbol{m}_{\text{des}}$, using (10) with $\boldsymbol{f}_{\text{des}} = \boldsymbol{0}$. However, we have already identified that at least one of our actuators is saturated, and we must determine if $\boldsymbol{\tau}_{\boldsymbol{m}}$ is achievable (given that some of our actuation is being used to apply $\gamma\boldsymbol{\tau}_{\boldsymbol{f}}$), and if not, if some scaled-down version of $\boldsymbol{m}_{\text{des}}$ is achievable. To check if the moment is achievable, we apply a $-\gamma\boldsymbol{\tau}_f$ offset to the upper and lower saturation limits on the joint torques. We then verify that $\boldsymbol{\tau}_{\boldsymbol{m}}$ is achievable, and if not, we linearly scale down the requested torque by a constant $0 < \delta < 1$ to ensure that none of that of the actuators is beyond saturation.

The complete pseudocode to implement the hybrid force-moment braking pulse is provided as Algorithm 1.

V. Human-Subjects Study

We conducted a humans-subjects study to test the hypothesis that hybrid force-moment braking pulses result in haptic virtual surfaces that are perceived as harder than surfaces using traditional braking pulses, without harming the realism of the surface. The study consists of two experiments that are identical in every way except in the question that is posed to the subject, with each utilizing a distinct set of human subjects. In the first experiment, the subjects are asked to determine which surface feels harder. In the second experiment, the subjects are told that the virtual surface is designed to simulate a hard metal surface, and the subject is asked to determine which surface feels more realistic.

A. Subjects

The first experiment is performed by 10 (5 male, 5 female) subjects, whose ages range from 22 to 28 years, and whose weights range from 60 to 95 kg. The second experiment is performed by a different 10 (5 male, 5 female) subjects, whose ages range from 21 to 35 years, and whose weights range from 65 to 95 kg. Subjects are all right-handed, and have normal tactile sensation and normal (corrected) vision, by self-report. The study was approved by the University of Utah Institutional Review Board.

B. Apparatus

The experimental setup for our study is shown in Fig. 2. Experiments were conducted with an Entact W6D haptic device, which comprises closed-chain kinematics with seven backdrivable motors. The complete kinematic equations are provided in Appendix I. The maximum joint torques for each of the six main (most proximal) motors is $\|\tau_{\max}\| = 511\,\text{N·mm}$ (from the device's header file). Our algorithm does not make use of the small distal motor responsible for torque about the axis of the stylus (although its maximum value is $\|\tau_{\max}\| = 145\,\text{N·mm}$). Because we are not certain what factor is ultimately causing the enforced maximum motor-torque value—which could be either amplifier saturation, power-supply saturation, or motor safety limits—we choose to artificially saturate the motors at $\|\tau_{\max}\| = 255\,\text{N·mm}$, such that our investigation of the underlying force-moment illusion will not be confounded by the unknown parameters of the commercial haptic device.

A simple haptic virtual surface in a horizontal plane was implemented using OpenGL and freeGLUT libraries. The surface was designed with a relatively small area (see the monitor in Fig. 2), which ensured that the subject's arm, wrist, and hand were in approximately the same posture throughout the experiment. This nominal configuration is defined by the haptic-device joint angles $\theta_1 = -36.4\,\text{rad}$, $\theta_2 = 81.1\,\text{rad}$, $\theta_3 = 9.44\,\text{rad}$, $\theta_4 = -48.0\,\text{rad}$, $\theta_5 = 87.0\,\text{rad}$, $\theta_6 = 6.00\,\text{rad}$, and θ_7 undetermined; see Appendix I for more details. At this nominal configuration, the maximum upward force that can be rendered by the device before saturation is 4.90 N, and with our artificially imposed motor saturations the maximum upward force is 2.45 N It was determined in pilot testing that a relatively high stiffness value (i.e., near the maximum value) that could be rendered without any perceptible unstable/nonpassive artifacts was 5678 N/m. This stiffness value was used consistently throughout all experiments.

We calculate the desired (traditional) braking pulse using (1), with $N = 1$ and $T = 0.001\,\text{s}$. We calculate the incoming momentum of the hand using $\boldsymbol{p} = m\boldsymbol{v}$. We estimated the mass of the subject's hand as 0.6% of their total body weight [19]. We calculate the velocity of the HIP with (7), using the joint-velocity estimates provided by the Entact W6D SDK.

To ensure that the braking pulses were discrete contact events, after a contact is detected and a braking pulse is implemented, the HIP must be raised 3 mm above the haptic virtual surface before another braking pulse can be given.

Fig. 2: The experimental setup. The subject holds the haptic device's stylus, with his elbow resting on an armrest, and controls a proxy in a simple 3D virtual world. A dot was painted on the stylus, indicating where the subject should grasp it, with $\|\boldsymbol{r}\| = 120\,\text{mm}$. The user can use the keyboard to toggle between the surface properties and enter the final answer.

We found in pilot testing that this was sufficient to allow free tapping of the surface without any noticeable artifacts.

C. Design

In each experiment, the subject compares two surfaces: one rendered using a traditional force braking pulse, and one using a hybrid force-moment braking pulse. Each subject compares $n = 20$ such pairs. Each type of braking pulse is presented first in half of the trials and second in half of the trials, with the presentations randomly sorted. To evaluate statistical significance, we use the binomial distribution, which states that the probability P of k selections in n trials if the differences are chosen by random chance ($p = 0.5$) is calculated as:

$$P = \binom{n}{k} p^k (1-p)^{n-k} = \frac{n!}{k!(n-k)!} p^k (1-p)^{n-k} \qquad (12)$$

Considering the cumulative probability centered at $k = 10$, using the conventional significance of $\alpha = 0.05$, a subject must select a given braking-pulse method at least 15 times out of the 20 trials to conclude that there is a significant difference from random chance. As such, $k \geq 15$ indicates a significant preference for the hybrid method, whereas $k \leq 5$ indicates a significant preference for the traditional method.

In each experiment, the method described above is used to classify each subject as a binary success or failure, where "success" is defined as a subject for which $k \geq 15$, and "failure" is defined by $k < 15$. We then apply the Wilson score interval (see Appendix II) to our results with 10 human subjects to determine the 95% confidence intervals for our results that are expected in the broader population.

D. Procedure

The subject is seated at a desk, in front of a computer monitor and keyboard, with the haptic device placed to the

right of the monitor. The subject is instructed to adjust the chair and armrest until they can comfortably hold the stylus such that their virtual proxy is hovering near the virtual surface on the monitor. The subject is instructed to grasp the stylus such that the red band on the stylus sits at the midpoint between the two points of contact, which are the grasp point between the fingers and the purlicue (i.e., the crook of the thumb). The subject is instructed to hold the stylus pointing down and away from them, with $\phi \approx 45°$, and is shown by the experimenter what this means (see Fig. 2). Underneath the virtual surface, a "1" or "2" appears on the monitor. Subjects are instructed to tap on the surface, pressing the "T" key on the keyboard to toggle between surface 1 and surface 2, until they determine which surface feels harder or more realistic of a hard metal surface (depending on the experiment), at which point they press the "1" or "2" key to enter their choice. The trial number displayed at the top of the monitor updates to give a visual confirmation that the keypress was detected. A new choice cannot be entered until the "T" key has been pressed again, which eliminates the possibility of accidentally entering an answer twice, or of entering an answer before both surfaces have been felt in a given trial. Subjects are told that tapping harder or faster may make it easier to distinguish differences between the surfaces. Subjects are instructed that if it is impossible to distinguish between the two surfaces then they should randomly select "1" or "2". Subjects are told that they may take a break at any time during the experiment. Throughout the study, the total completion times across subjects was 15–25 min.

E. Results

1) Experiment 1: In the experiment in which we asked the subjects which surface felt harder (Fig. 3(a)), nine out of ten subjects chose the hybrid force-moment braking pulse at least 15 times out of 20 trials, meaning that they perceived those surfaces as harder (with at least 95% confidence). The remaining subject chose the hybrid force-moment braking pulse 12 out of 20 trials, which is not enough to be considered significantly different than random chance.

With nine out of ten subjects indicating that the hybrid force-moment braking pulse feels harder, we can place a 95% confidence interval on what we should expect in the population more broadly at 70–98% of the population. Thus, we can conclude that a supermajority of people will perceive the illusion of a harder surface.

2) Experiment 2: In the experiment in which we asked the subjects which surface felt more realistic as a simulation of a hard metallic surface (Fig. 3(b)), seven out of ten subjects chose the hybrid force-moment braking pulse at least 15 times out of 20 trials, meaning that they perceived those surfaces as more realistic (with at least 95% confidence). Two subjects chose the hybrid force-moment braking pulse in the majority of trials, but not enough to be considered significantly different than random chance. The remaining subject never selected the hybrid force-moment braking pulse.

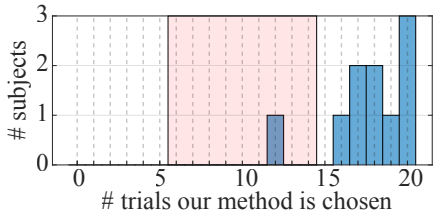

(a) Experiment 1: "Which surface felt harder?"

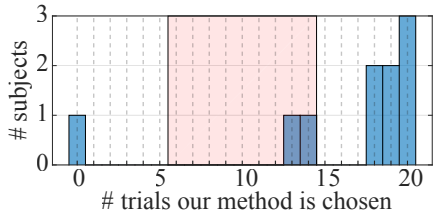

(b) Experiment 2: "Which surface was a more realistic simulation of a hard metallic surface?"

Fig. 3: Histogram of subjects' responses for each experiment. The pink areas indicate values that are not different from random chance (using a significance $\alpha = 0.05$).

With one out of ten subjects indicating that the hybrid force-moment braking pulse feels less realistic of a hard metal surface, we can place a 95% confidence interval on what we should expect in the population more broadly at 2–40% of the population. Thus, we can conclude that a minority of people will perceive that the surface feels less realistic.

VI. Discussion

We have shown that is it possible to use a hybrid force-moment braking pulse, which makes use of a newly identified haptic illusion based on a nonrealistic applied moment, to create virtual surfaces that feel harder than those rendered with traditional braking pulses, without harming the realism of the surface, for the majority of users. This new haptic illusion can be used to increase the performance of existing haptic interfaces, seemingly pushing them beyond their saturation limits. Although we only incorporated the haptic illusion with a basic braking pulse here, it could easily be incorporated with other more sophisticated event-based haptic feedback.

During our initial investigation, we wondered whether the illusion that we were feeling was due to the increase in the reaction moment at the hand as hypothesized, or if was simply due to the introduction of additional high-frequency haptic or audio stimulus. To test this alternate hypothesis, we reversed the sign on the commanded moment $\boldsymbol{m}$. For example, one can imagine the case of Fig. 1(b) if the sign of τ_5 was reversed. What we experienced was a perceived softening of the virtual surface, which led us to dismiss the alternate hypothesis.

In our study, we only considered a single haptic device, at a single nominal configuration, at a single arm-wrist-hand posture, with a single braking-pulse duration of 0.001 s, with a single stiffness of the underlying virtual surface, and with a

single method to estimate the momentum of the hand during contact. As a result, although we have demonstrated the existence of a new haptic illusion, characterizing the bounds of when this haptic illusion can be exploited is left as an open problem. As discussed earlier, we know the value of ϕ will certainly make a difference. The specific haptic device and its configuration are both likely to make a difference as well, as is the relative maximum strength of the device's actuators. It is also likely that extending the duration of the braking pulse too long, in an attempt to increase the magnitude of the pulse, will likely harm the realism of the sensation. It is unclear how a change in the stiffness of the underlying virtual surface would affect the illusion.

In the experiment in which we asked the subjects which surface felt more realistic as a simulation of a hard metallic surface, one of the subjects never chose the surface with the hybrid force-moment braking pulse. That subject told the experimenter that her choice was because the surface with the hybrid force-moment braking pulse felt "strange". It is unclear to us what caused this to happen. In the experimenter's perception, this subject was not holding the stylus differently or interacting with the virtual surface differently than the other subjects. This suggests that the haptic illusion is not effective for everybody, but our statistical analysis indicates that the majority of people will not exhibit this problem with the haptic illusion. It is also worth noting that the robustness of this haptic illusion is comparable to many other well known haptic illusions [9]. That being said, when incorporating this haptic illusion into virtual environments, it will likely be best to incorporate the ability to toggle the illusion on-off for those who feel it harms the realism of the virtual environment.

APPENDIX I
KINEMATICS OF THE ENTACT W6D

We are interested in the velocity $\boldsymbol{v} = \begin{bmatrix} v_x & v_y & v_z \end{bmatrix}^\top$ and angular velocity $\boldsymbol{\omega} = \begin{bmatrix} \omega_x & \omega_y & \omega_z \end{bmatrix}^\top$ at the HIP (see Fig. 4). These vectors can be expressed in either the workspace coordinate frame "w" or the stylus coordinate frame "s". The representations in these two frames are related by a rotation matrix ${}^w\mathbb{R}_s = {}^s\mathbb{R}_w^\top$ describing the orientation of the stylus frame with respect to the workspace frame:

$$ {}^w\boldsymbol{v} = {}^w\mathbb{R}_s\, {}^s\boldsymbol{v}, \quad {}^w\boldsymbol{\omega} = {}^w\mathbb{R}_s\, {}^s\boldsymbol{\omega} \tag{13} $$

The Entact W6D haptic device comprises two parallel arms (Arms 0 and 1), each with three degrees of freedom, and a distal motor (M7) that is responsible for torque about the axis of the stylus. The velocity and angular velocity at the HIP can be mapped to the velocities $\boldsymbol{v}_0$ and $\boldsymbol{v}_1$ at the ends of Arms 0 and 1, respectively, and the angular velocity about the axis of the stylus is treated independently:

$$ \begin{bmatrix} {}^sv_{0x} \\ {}^sv_{0y} \\ {}^sv_{0z} \\ {}^sv_{1x} \\ {}^sv_{1y} \\ {}^sv_{1z} \\ {}^s\omega_z \end{bmatrix} = \underbrace{\begin{bmatrix} 1 & 0 & 0 & 0 & L & 0 \\ 0 & 1 & 0 & -L & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}}_{\mathbb{A}} \begin{bmatrix} {}^sv_x \\ {}^sv_y \\ {}^sv_z \\ {}^s\omega_x \\ {}^s\omega_y \\ {}^s\omega_z \end{bmatrix} \tag{14} $$

Each of the two arms, denoted by Arm i with $i \in \{0,1\}$, has a simple kinematic structure, with its end position in the workspace frame described by

$$ {}^w\boldsymbol{d}_i = \begin{bmatrix} {}^wx_i & {}^wy_i & {}^wz_i \end{bmatrix}^\top \tag{15} $$

where

$$ \begin{aligned} {}^wx_i =& (2i-1)l_0 \sin(\theta_{3i+3}) - d_1 + l_2 \cos(\theta_{3i+1}) \cos(\theta_{3i+3}) \\ &+ \underbrace{l_3 \cos(\theta_{3i+3}) \cos(\theta_{3i+1} + \theta_{3i+2})}_{{}^w\boldsymbol{a}_i[1]} \end{aligned} \tag{16} $$

$$ {}^wy_i = -l_1 + l_2 \sin(\theta_{3i+1}) + \underbrace{l_3 \sin(\theta_{3i+1} + \theta_{3i+2})}_{{}^w\boldsymbol{a}_i[2]} \tag{17} $$

$$ \begin{aligned} {}^wz_i =& (2i-1)l_0 \cos(\theta_{3i+3}) + d_2 - l_2 \cos(\theta_{3i+1}) \sin(\theta_{3i+3}) \\ &- (2i-1)d_3 \underbrace{-l_3 \sin(\theta_{3i+3}) \cos(\theta_{3i+1} + \theta_{3i+2})}_{{}^w\boldsymbol{a}_i[3]} \end{aligned} \tag{18} $$

with the orientation of the stylus frame described by

$$ {}^w\mathbb{R}_s = \begin{bmatrix} {}^w\hat{\boldsymbol{x}}_s & {}^w\hat{\boldsymbol{y}}_s & {}^w\hat{\boldsymbol{z}}_s \end{bmatrix} \tag{19} $$

where

$$ {}^w\hat{\boldsymbol{z}}_s = \frac{{}^w\boldsymbol{d}_0 - {}^w\boldsymbol{d}_1}{\| {}^w\boldsymbol{d}_0 - {}^w\boldsymbol{d}_1 \|} \tag{20} $$

$$ {}^w\hat{\boldsymbol{x}}_s\Big|_{\theta_7=0} = \frac{{}^w\boldsymbol{a}_1 \times {}^w\hat{\boldsymbol{z}}_s}{\| {}^w\boldsymbol{a}_1 \times {}^w\hat{\boldsymbol{z}}_s \|} \tag{21} $$

$$ {}^w\hat{\boldsymbol{y}}_s\Big|_{\theta_7=0} = {}^w\hat{\boldsymbol{z}}_s \times \left({}^w\hat{\boldsymbol{x}}_s\Big|_{\theta_7=0} \right) \tag{22} $$

$$ \begin{bmatrix} {}^w\hat{\boldsymbol{x}}_s & {}^w\hat{\boldsymbol{y}}_s \end{bmatrix} = \begin{bmatrix} \cos(\alpha\theta_7) & -\sin(\alpha\theta_7) \\ \sin(\alpha\theta_7) & \cos(\alpha\theta_7) \end{bmatrix} \begin{bmatrix} {}^w\hat{\boldsymbol{x}}_s\Big|_{\theta_7=0} & {}^w\hat{\boldsymbol{y}}_s\Big|_{\theta_7=0} \end{bmatrix} \tag{23} $$

where $\alpha = {}^s\omega_z / \dot{\theta}_7 = 0.185$ is the transmission ratio due to the gearhead on M7 [20].

The Jacobians of the forward kinematics of the arms, $\mathbb{J}_{0\text{fk}}$ and $\mathbb{J}_{1\text{fk}}$, are found by differentiating the forward kinematics,

$$ \mathbb{J}_{i\text{fk}} = \begin{bmatrix} \frac{\partial {}^w\boldsymbol{d}_i}{\partial \theta_1} & \cdots & \frac{\partial {}^w\boldsymbol{d}_i}{\partial \theta_7} \end{bmatrix}, \tag{24} $$

which can be calculated numerically using central differences.

TABLE I: Measured lengths (mm) from Entact W6D.

l_0	l_1	l_2	l_3	d_1	d_2	d_3
42.0	34.0	180	165	46.0	205	69.7

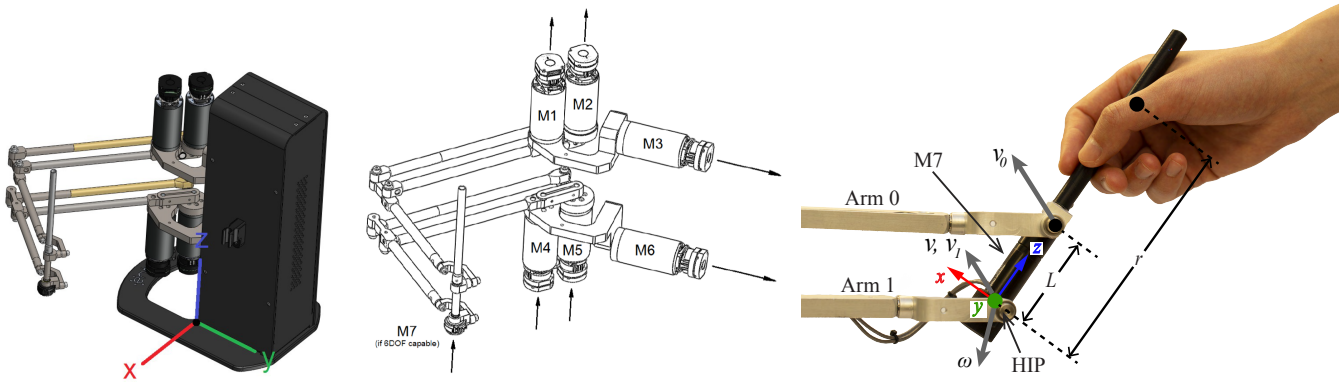

(a) (b) (c)

Fig. 4: Entact W6D. (a) Definition of workspace frame. (b) Kinematic structure. (c) Close-up of the stylus and hand, with parameters used in this study. The origin of the stylus coordinate frame is located at the HIP. The stylus coordinate frame is depicted for $\theta_7 = 0$. $\boldsymbol{v}$ and $\boldsymbol{\omega}$ are the velocity and angular velocity, respectively, of the stylus at the HIP. We defined the upper arm as "Arm 0" and the lower arm as "Arm 1". Since the seventh (most distal) motor is mounted on the stylus, and it is decomposed from the rest of the device, we treat it as an independent actuator. The HIP and the end of Arm 1 are collocated, resulting in $\boldsymbol{v} = \boldsymbol{v}_1$. Images (a) and (b) are reprinted from [20] with permission of Ryan Leslie, Entact Robotics.

The inverses of these matrices enable us to compute the joint velocities from $\boldsymbol{v}_0$, $\boldsymbol{v}_1$, and ${}^s\omega_z$:

$$\begin{bmatrix} \dot{\theta}_1 \\ \dot{\theta}_2 \\ \dot{\theta}_3 \\ \dot{\theta}_4 \\ \dot{\theta}_5 \\ \dot{\theta}_6 \\ \dot{\theta}_7 \end{bmatrix} = \underbrace{\begin{bmatrix} \mathbb{J}_{0\text{fk}}^{-1} & 0 & 0 \\ 0 & \mathbb{J}_{1\text{fk}}^{-1} & 0 \\ 0 & 0 & \frac{1}{\alpha} \end{bmatrix}}_{\mathbb{B}} \begin{bmatrix} {}^w v_{0x} \\ {}^w v_{0y} \\ {}^w v_{0z} \\ {}^w v_{1x} \\ {}^w v_{1y} \\ {}^w v_{1z} \\ {}^s \omega_z \end{bmatrix} \tag{25}$$

Combining the preceding results, we can calculate the joint velocities for a given stylus velocity and stylus angular velocity in a form that is equivalent to (8):

$$\begin{bmatrix} \dot{\theta}_1 \\ \dot{\theta}_2 \\ \dot{\theta}_3 \\ \dot{\theta}_4 \\ \dot{\theta}_5 \\ \dot{\theta}_6 \\ \dot{\theta}_7 \end{bmatrix} = \underbrace{\mathbb{B} \begin{bmatrix} {}^w\mathbb{R}_s & 0 & 0 \\ 0 & {}^w\mathbb{R}_s & 0 \\ 0 & 0 & 1 \end{bmatrix} \mathbb{A} \begin{bmatrix} {}^s\mathbb{R}_w & 0 \\ 0 & {}^s\mathbb{R}_w \end{bmatrix}}_{\mathbb{J}_{\text{ik}}} \begin{bmatrix} {}^w v_x \\ {}^w v_y \\ {}^w v_z \\ {}^w \omega_x \\ {}^w \omega_y \\ {}^w \omega_z \end{bmatrix} \tag{26}$$

Appendix II
Wilson Score Interval

For an estimated probability of $\hat{p}$ from n trials, a conventional confidence interval, which assumes a normal distribution centered at $\hat{p}$, is calculated as

$$\hat{p} \pm z \sqrt{\frac{\hat{p}(1-\hat{p})}{n}} \tag{27}$$

where $z = 1.96$ for a 95% confidence interval (i.e., for a significance of $\alpha = 0.05$). In some cases, this can result in confidence intervals that stretch below 0 or above 1, which is not meaningful for probability.

The Wilson score interval [21] is an improvement to the conventional method of calculating confidence intervals when using binomial distributions. It provides better estimates of the true interval, particularly in cases with a small number of trials or with extreme probability values. The Wilson score interval is calculated as

$$\frac{\hat{p} + \frac{z^2}{2n}}{1 + \frac{z^2}{n}} \pm \frac{z}{1 + \frac{z^2}{n}} \sqrt{\frac{\hat{p}(1-\hat{p})}{n} + \frac{z^2}{4n^2}} \tag{28}$$

where the same z value is used as above. Note that this interval is not centered on $\hat{p}$.

References

[1] J. E. Colgate and G. Schenkel, "Passivity of a class of sampled-data systems: Application to haptic interfaces," *J. Robot. Syst.*, vol. 14, no. 2, p. 3747, 1997.

[2] J. J. Abbott and A. M. Okamura, "Effects of position quantization and sampling rate on virtual-wall passivity," *IEEE Trans. Robot.*, vol. 21, no. 5, p. 952964, 2005.

[3] N. Diolaiti, G. Niemeyer, F. Barbagli, and J. K. Salisbury, Jr., "Stability of haptic rendering: Discretization, quantization, time-delay and coulomb effects," *IEEE Trans. Robot.*, vol. 22, no. 2, p. 256268, 2006.

[4] D. A. Lawrence, L. Y. Pao, A. M. Dougherty, M. A. Salada, and Y. Pavlou, "Rate-hardness: A new performance metric for haptic interfaces," *IEEE Trans. Robot. Autom.*, vol. 16, no. 4, pp. 357–371, 2000.

[5] G. Han and S. Choi, "Extended rate-hardness: A measure for perceived hardness," in *EuroHaptics 2010, Part I, LNCS 6191*, A. M. L. Kappers, J. B. F. van Erp, W. M. Bergmann Tiest, and F. C. T. van der Helm, Eds. Springer-Verlag Berlin Heidelberg, 2010, pp. 117–124.

[6] S. E. Salcudean and T. Vlaar, "On the emulation of stiff walls and static friction with a magnetically levitated input/output device," *J. Dyn. Syst.-T. ASME*, vol. 119, pp. 127–132, 1997.

[7] A. M. Okamura, M. Cutkosky, and J. Dennerlein, "Reality-based models for vibration feedback in virtual environments," *IEEE/ASME Trans. Mechatronics*, vol. 6, no. 3, pp. 245–252, 2001.

[8] K. J. Kuchenbecker, J. Fiene, and G. Niemeyer, "Improving contact realism through event-based haptic feedback," *IEEE Trans. Vis. Comput. Graph.*, vol. 12, no. 2, pp. 219–230, 2006.

[9] S. J. Lederman and L. A. Jones, "Tactile and haptic illusions," *IEEE Trans. Haptics*, vol. 4, no. 2, pp. 273–294, 2011.

[10] H. Yano, M. Yoshie, and H. Iwata, "Development of a non-grounded haptic interface using the gyro effect," *Proc. Haptics Symp.*, pp. 32–39, 2003.

[11] K. N. Winfree, J. M. Romano, J. Gewirtz, and K. J. Kuchenbecker, "Control of a high fidelity ungrounded torque feedback device: The iTorqU 2.1," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2010, pp. 1347–1352.

[12] J. M. Walker, H. Culbertson, M. Raitor, and A. M. Okamura, "Haptic orientation guidance using two parallel double-gimbal control moment gyroscopes," *IEEE Trans. Haptics*, vol. 11, no. 2, p. 267278, 2018.

[13] S. A. Pedram, R. L. Klatzky, and P. Berkelman, "Torque contribution to haptic rendering of virtual textures," *IEEE Trans. Haptics*, vol. 10, no. 4, pp. 567–579, 2017.

[14] R. Zhang, A. J. Boyles, and J. J. Abbott, "Six principal modes of vibrotactile display via stylus," in *Proc. IEEE Haptics Symp.*, 2018, pp. 313–318.

[15] R. Zhang, T. J. Schwer, and J. J. Abbott, "Dimensional reduction for 6D vibrotactile display," *IEEE Trans. Haptics*, vol. 13, no. 1, pp. 102–108, 2020.

[16] J. Park, Y. Oh, and H. Z. Tan, "Effect of cutaneous feedback on the perceived hardness of a virtual object," *IEEE Trans. Haptics*, vol. 11, no. 4, pp. 518–530, 2018.

[17] K. M. Lynch and F. C. Park, *Modern Robotics: Mechanics, Planning, and Control*. Cambridge University Press, 2017.

[18] V. Kumar and J. F. Gardner, "Kinematics of redundantly actuated closed chains," *IEEE Trans. Robot. Autom.*, vol. 6, no. 2, pp. 269–274, 1990.

[19] D. A. Winter, *Biomechanics and Motor Control of Human Movement*. John Wiley and Sons, 1990.

[20] *W5D User Manual*, Entact Robotics.

[21] E. B. Wilson, "Probable inference, the law of succession, and statistical inference," *J. Am. Stat. Assoc.*, vol. 22, no. 158, pp. 209–212, 1927.