THE “HAPTIC CAMERA”: A TECHNIQUE FOR CHARACTERIZING AND PLAYING BACK HAPTIC PROPERTIES OF REAL ENVIRONMENTS

Karon E. MacLean
Interval Research Corporation
1801 Page Mill Road
Palo Alto, CA 940304, USA
maclean@interval.com

ABSTRACT

In some applications of simulated haptic feedback, it is necessary not only that the display feel stable and adequately realistic: it must resemble with high fidelity a particular real device, emulating the feedback supplied by a physical environment that does or will exist. An automated haptic characterization technique was developed in which the haptic display itself is used under tight position control to probe the real environment while measuring interaction forces. A model structure is assumed which may be nonlinear and globally discontinuous but piecewise continuous, and is refined for successively higher-order parameters by application of appropriately chosen position trajectories. A one degree-of-freedom linear-acting motor with 90 mm stroke, a peak force of 60 N and equipped with position and force sensors was used as both a probe and a force display.

This algorithm was tested on one real environment, a toggle switch with a nonlinear, compliance-dominated impedance. Because of the nature of the real device, only stiffnesses were extracted; this zero-order, piecewise linear model assumption provided a close fit to the measured trajectory with under 30 seconds of measurement and processing time. This was a substantial improvement in playback fidelity over that achieved by other characterization methods. Future work includes testing the algorithm on higher order systems.

INTRODUCTION

Creating virtual haptic feedback usually requires a parameterized model of a physical device or environment. In some cases the virtual environment is intended to resemble (target) another that actually does or might in the future exist: examples include a virtual prototype of a real device to be physically constructed, or a surgeon or aviator training system whose utility depends on how well it captures key features of the real environments to which its students will graduate. In these cases, the model must correspond in some manner and degree to the real environment. Creating, parameterizing and stabilizing this model to an acceptable degree of fidelity is often difficult, and success depends on the model, the nature of the targeted environment and the suitability of the display hardware and control. This paper focuses on the challenge of characterizing a real environment for the purpose of haptic emulation.

Modeling for Feeling

An emulation's fidelity is related to the quality of modeling and parameterization of its real target, but the definition of a good model depends on how the model will be used. Modeling tools are available to predict behavior of a mechanical system under specified excitation conditions, and model adequacy is understood. In this case, however, the task is to create an input which triggers the same sensations as does interaction with the emulation's target. Those in psychophysics are developing models for the mechanisms of the haptic sense (Dandekar, 1995; Pang, 1991); but at this time we do not have precise knowledge of what we feel, far less how. In order to begin, here we assume that for the purpose of creating a haptic emulation, a human's haptic perception coincides with what an engineer sees when he/she examines a mechanical system and approximates its behavior with interconnecting masses, springs, dampers and other more complex elements. It is important to recognize that this is an incomplete picture: we feel through a filter of our sensory system and finger, hand and arm impedance (Gotow, 1989; Milner and Franklin, 1995) and the intrinsic impedance of the force display itself (Jones and Hunter, 1990). Perception is influenced by exploration strategy as well as our sensing apparatus (Lederman and Klatzky, 1993).

Simple environments, adequately described by a single phenomenon or even a single parameter, can also be difficult to model and parameterize. A viscosity-dominated slide potentiometer might be completely dissipative, but pure viscous damping is rarely encountered outside of engineering scratchpads, and need not be linear.

Another modeling complication is posed by the complex elements such as higher order dynamics which engineers often neglect in order to make a problem tractable and when these elements do not contribute to gross features of the system's behavior. Difficult to model and parameterize, such elements are frequently responsible for a critical aspect of a person's haptic perception. When this simplified model is used for a haptic emulation, the high order dynamics are usually hard to play back because most haptic display hardware is geared towards the frequency range of human voluntary motion (0-10 Hz); achieving both low and high frequency actuation generally will require tandem actuation (Kontarin and Howe, 1995; Morrell and Salisbury, 1995; Pratt and Williamson, 1995). The emulation of high-frequency nonlinear phenomena such as wall-tapping and stiff detents demand high bandwidth hardware and occasionally sophisticated control.

METHODS OF IDENTIFICATION

There are a variety of ways in which a real environment may be characterized, each with its own tradeoffs. Discussed below, they include (I) manual estimation, i.e. adjusting the emulation until it "feels right"; and (II) explicit measurement of structural elements of a physical model, e.g. weighing of a mass. The first suffers from a lack of repeatability and objectivity and often proves non-convergent in the characterization of complicated mechanisms; the second may entail difficult measurements. The active probe concept (III) is the basis of the technique described in the remainder of this paper. In this research, the best results were produced by combining methods I, II and III; this strategy is Method IV.

I. "Guess": Iterative Perceptual Matching: A model is assumed, not necessarily based on the actual mechanism, with parameters which approximate the perceived force profile. The emulation designer iteratively closes the gap perceived between the emulation and the real target by adjusting the model and its parameters. This approach is low overhead, requiring only a crude concept of the mechanism. Moreover, inexact as it might seem, this mode of measurement which relies on the designer's own haptic perception is the "holy grail" of emulation fidelity. When other modeling approaches deliver slightly off-target results, iterative manual matching are inevitably be used to fine tune them.

Disadvantages are many. The reliance on haptic perception is good for fidelity but poor from the concern of a quantitative mapping between real and virtual environments, e.g. for building a haptic prototype. Dependent on exploration strategy, haptic perception is not particularly accurate, objective nor repeatable as a measurement tool. Because it is iterative, the method is time consuming. If the model is complex or not physically based, the connection between parameterizations and output feel might be indirect and nonintuitive, and the modeling process nonconvergent. A model constructed in this manner may be simulated mathematically to verify response, but if it does not closely resemble the real mechanism this might not lend insight into how the model could be changed to improve fidelity.

II. Structural Modeling: The real device's mechanism is modeled and parameterized through disassembly and measurement of individual components. For a bat-type spring toggle switch of the sort found on laboratory equipment, this would mean disassembling the toggle switch body and measuring the spring constant for the spring in the actuator, the friction coefficient of the actuator tip sliding on the internal cantilever, and all geometric dimensions. This model may also be mathematically simulated to see if the predicted behavior matches that of the real target, probably with greater benefit than in the previous method since the model can be compared with the real device for errors; e.g. unmodeled dynamics could be identified and added to the model. Correct behavior of the mathematically modeled mechanism alone is no proof that the model will be stable when run in a virtual environment controller and displayed on a system with an impedance of its own; inclusion of the emulator plant in the math simulation helps with this. More crucial is difficulty in ascertaining when all key perceptual components of this mechanically derived model have been included.

III. Active Probe: Impedance components of the fully assembled real device are characterized using a computer controlled force probe. The "Haptic Camera" is one algorithm for this approach, where an input is supplied to the system while its output is measured; in the current algorithm, a presupposition of model structure must be made. The system identification technique has the advantages of being objective, relying in its parameterization upon mechanical measures; and contextual in that it measures impedance empirically at the same location that a person manipulating the device will feel it, at the handle. In a successful implementation applied to an amenable identification problem, it is rapid. The method developed here completes data collection and analysis for one parameter identification in about 20 seconds.

Its disadvantages are in its higher overhead of hardware and program development, and in its assumption of model structure. A more powerful incarnation might be a stochastic identification which makes no such assumptions about structure.

Other algorithms could be applied to the "active probe" concept. A likely approach is to excite the system with a single frequency-rich stochastic signal rather than a sequence of narrow bandwidth trajectories (Liu and Asada, 1992). For some environments, this would be an appropriate and possibly efficient method. For systems distinguished by small geometries and narrow continuous regions separated by discontinuities in model and/or parameters, there are practical obstacles to imposing trajectories which repeatedly cross the discontinuities. It will generally be difficult to maintain tight position control across a severe discontinuity and through rapid direction and velocity changes, so these are points of greatest error. When it is desirable to locate the discontinuities but characterize model and parameters in the piecewise continuous zones, minimizing zone-crossings and direction changes helps to reduce characterization error. Since these are the sort of devices of most interest to this research, the separate narrow-bandwidth trajectory approach was taken here.

IV. All of the Above: In practice, these methods are best used in combination. Some knowledge of the underlying structural model is helpful as a starting point for an iterative perceptual approach (I), particularly if it is a simple structure and/or one whose primary features are easily captured in a model. Likewise, model-based approaches (II and III) will need tuning in any but the simplest case, because the model will be imperfect and emulator dynamics will influence the model output.

THE HAPTIC CAMERA

Algorithm

The identification algorithm is based on partitioning real environment impedance into its spatial and dynamic components, given an assumption of general model structure. Model parameterizations are extracted successively from the real environment, low order to high order, by applying tightly controlled position inputs of the appropriate type and then measuring and analyzing interaction force.

The probe in this implementation was a single degree of freedom linear-acting emulation system described in

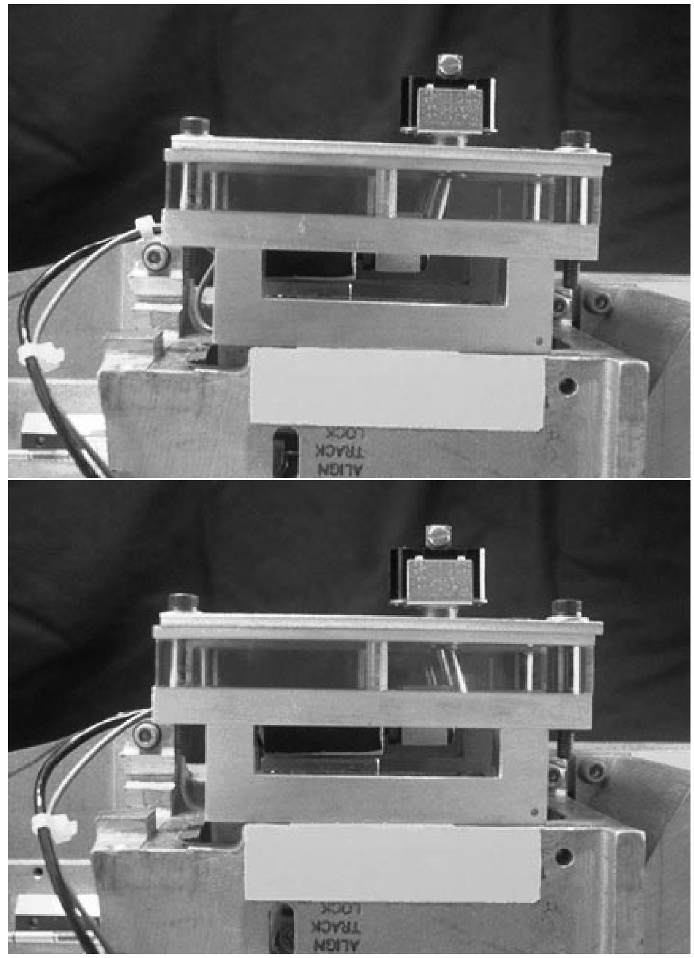

Figure 1: A force display used as a force probe. The target, a toggle switch, is inverted over the force probe. The probe has moved from left to right between the two photographs, deflecting the toggle bat.

(MacLean and Durfee, 1995; MacLean, 1996), with a 90 mm range of motion, a maximum sustained output force of 60 N, stiffness of about 20 N/mm and an isometric force bandwidth of 100 Hz. The force display's handle was replaced by a short stiff probe for this project (Figure 1). The system's stiff direct-drive design made it ideal for the characterization task. The target device was mounted on a stiff plate which was in turn bolted rigidly to the emulator frame with the device handle upside down on a path which intersected that of the force probe.

The data collection and analysis procedure is most easily explained using as an example an environment whose dynamics are locally linear, but which contains discontinuities separating piece-wise linear regions with different parameters. To identify mass, damping, stiffness and equilibrium points ($\hat{M}(x)$, $\hat{B}(x)$, $\hat{K}(x)$ and $\hat{x}_o(x)$, respectively) as a function of region through a linear acting point of contact, trajectories are used successively as described in Table 1. The actual process of obtaining the model parameters is more involved than the vector division suggested here; the most important step is in determining the edges of model regions so that the piecewise continuous approximation may be made.

In the case of the velocity trajectory, it is helpful to use a variety of constant velocities in order to ascertain non-linearity in B. The same reasoning might also apply to estimation of M; however, there is not much room over the few inches or millimeters usually available to vary acceleration, and nonlinear inertia is a less common phenomenon for simple mechanisms.

Each trajectory can be duplicated to reduce the effects of measurement noise. Some types of measurement noise may be distinguished from unmodeled dynamics through repeated traversals, and reduced through averaging in the time/position domain or with spectral methods. Systematic measurement error which accrues through faulty hardware connections, etc, may give similar errors with repeated trials.

Implementation

Figure 2 demonstrates Step 1 (characterization of $\hat{K}(x)$) of the automatic identification algorithm on a stiffness-dominated toggle switch. The switch, visible inverted in Figure 1, is a garden-variety "momentary" toggle; when pushed it snaps to a local momentary equilibrium but springs back to its normal undeflected position when released.

The data collected for this characterization consists of the solid curve in (a), representing the measured force/position curve for this device ($f_m(x_m)$). The normal (open) position is at the left end of the plot. [1] Initial contact is seen in the steep positive force slope at around 1 mm; travel from there to peak force at 5 mm represents gradual deflection of the spring toggle. Just after 5 mm, the handle drops "over the top" into the metastable momentary position. Data collection was halted at around 7 mm of deflection, just before the hard back stop of the switch was struck. Data collected after that point is from a stationary probe.

Smoothing and acausal differencing (Dohrmann, 1988) produces the dashed and dotted curves in (a), which are $\frac{d(f_m)}{dx_m}$ and $\frac{d^2(f_m)}{dx_m{}^2}$, respectively. Region transitions are located based on super-threshold local maxima in the second force derivative; and finally, an estimate of stiffness ($\hat{K}(x)$) together with corresponding equilibrium points ($\hat{x}_o(x)$, one point for each region of approximated constant stiffness) is made by finding the average slope and x-intersect of $\frac{d(f_m)}{dx_m}$, respectively, for the region.

[1] The data was collected as a time sequence, $f_m(t_k)$. If the position control had been perfect, then $f_m(x_m)$ and $f_m(t_k)$ would be identical for a constant-velocity trajectory. Small divergences from this ideal were handled by sorting $f_m(t_k)$ according to ascending $x_{m_k}(t_k)$ and interpolating the sorted sequence to account for temporal non-uniformities.

Figure 2 (b) shows the functional result of this estimation procedure. The solid line repeats $f_m(x_m)$ from (a) as a reference; the piecewise constant-slope dotted line is $\hat{f}(x)$, estimated by running the x_m sequence through the estimated $\hat{K}(x)$ and $\hat{x}_o(x)$.

Control Over Region Resolution: The characterization of Figure 2 produced a relatively coarse division of six transitions and five piecewise linear segments. With this procedure it is easy to control the region resolution, by adjusting the threshold value of peak $\frac{d^2(f_m)}{dx_m{}^2}$ allowed to trigger a region transition (see following description of the transition detection algorithm). Figure 3 shows the same data set analyzed with a lower transition-detecting threshold; it produces 21 transitions and an even tighter match of $\hat{f}(x)$ to $f_m(x_m)$.

Use of the two models in emulation suggests that the lower resolution characterization is preferable, at least in this case. The small difference in output force as a function of measured position is well below what people can detect, and in general it is desirable to minimize the number of regions in an emulation. Region crossing tends to increase the sensation of emulation activeness because of temporal resolution, in the same way that virtual walls are more difficult to render passively than are continuous springs.

Auto-Transition Detection Algorithm: The transition detection procedure is the key to the success of this identification method. It chooses transitions not based on a threshold of arbitrarily chosen absolute magnitude, but on a relative difference approach which locates and retains natural groupings.

The transition detection algorithm can be summarized as follows:

1. Begin with second derivative of measured force. We know that
$$\frac{d^2(f_m)}{dx_m{}^2} = \frac{dK(x)}{dx_m}.$$
2. Locate all local maxima and minima in $\frac{d^2(f_m)}{dx_m{}^2}$, and sort in descending order of their amplitude's absolute value into a sequence $[\mathrm{amp}_k, \mathrm{amp}_{k-1}, \mathrm{amp}_{k-2}, \ldots]$ where amp_k is the largest in the sequence. Region transitions correspond to these peaks.

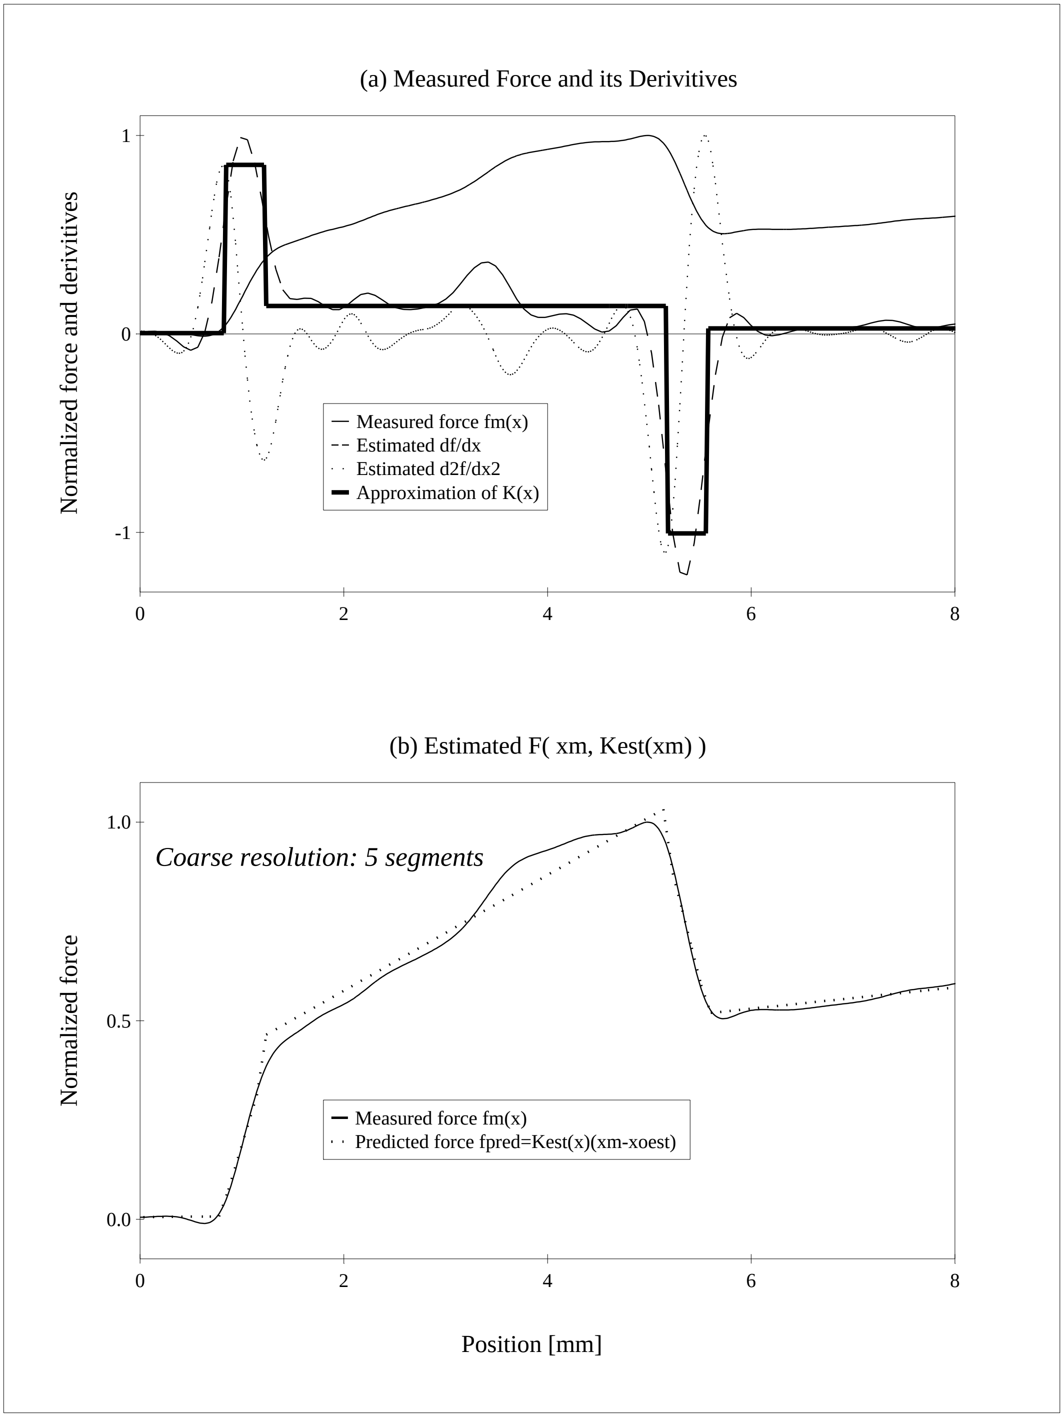

Figure 2: The identification algorithm on a stiffness-dominated system, finding zone transitions at a coarse resolution (6 transitions found). (a) Measured force as a function of controlled position ($f_m(x_m)$), and derived $\frac{d(f_m)}{dx_m}$ and $\frac{d^2(f_m)}{dx_m{}^2}$. The heavy segmented lines are regions of constant estimated $\hat{K}(x)$. (b) Measured force ($f_m(x_m)$) and predicted output force ($\hat{f}(x)$) based on estimated $\hat{K}(x)$ and $\hat{x}_o(x)$.

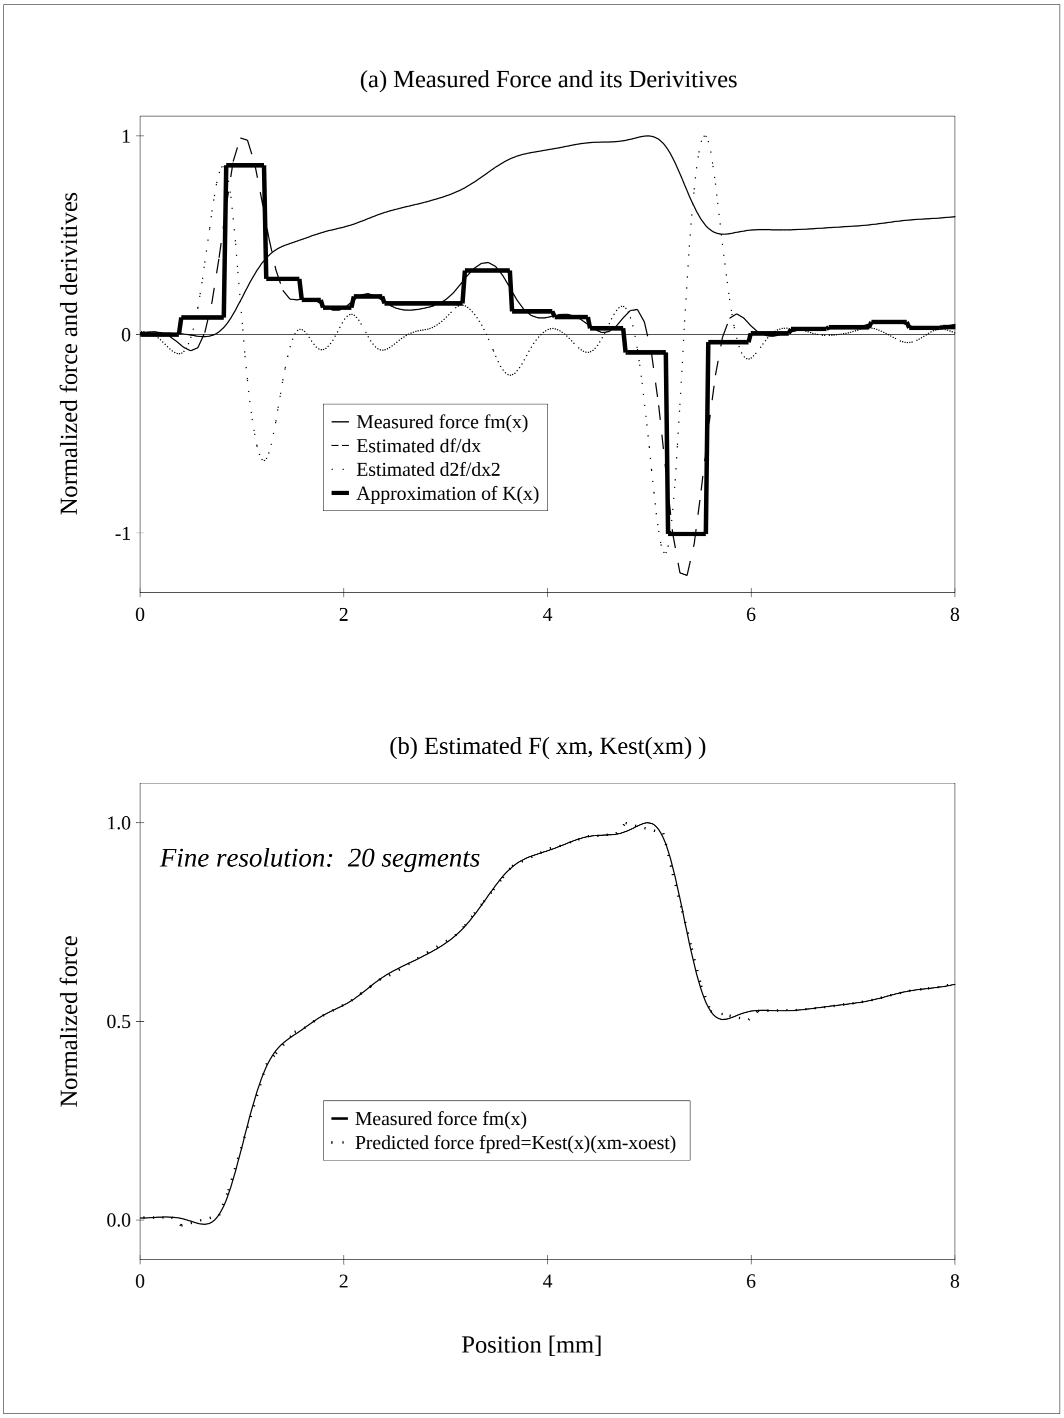

Figure 3: The same example and curves as in the previous example, but here set to find transitions at fine resolution (21 transitions or 20 segments found).

Computation Step	Trajectory Type	Information Extracted
1	quasi-static	$\hat{K}(x) = \frac{\vec{f}_m}{(\vec{x}_m - \hat{x}_o(x))}$
2	constant velocity	$\hat{B}(x) = \frac{\vec{f}_m - \hat{K}(x)\Delta x}{\vec{v}_m}$
3	constant acceleration	$\hat{M}(x) = \frac{\vec{f}_m - \hat{K}(x)\Delta x - \hat{B}(x)\vec{v}_m}{\vec{a}_m}$

Table 1: Successive trajectories used to identify model parameters

3. However, there are too many. Select the important ones by setting a cutoff value for

$$C = \frac{\text{amp}_{k-1}}{\text{amp}_k},$$

that is, the ratio of each amplitude to the next highest. This ratio locates transitions corresponding to divisions between natural groupings of amplitudes. A smaller value for C produces more regions.

In the example shown, the analysis which produced 6 regions used C=0.6, whereas the second analysis producing 20 regions used C=0.3.

Observations

At this time, the algorithm has been tested to the degree demonstrated: characterization of stiffness only, on a single real device. With this limited test set the results are promising. A completely automated assessment of $\hat{K}(x)$ and $\hat{x}_o(x)$, in playback with zero tuning felt closer to the real device than emulations produced previously by purely manual estimation. With minimal tuning (discontinuity handling, etc.; some of those modifications were standard treatment and could also be automated), it was quite recognizable as a simulation of the real device. [2] The procedure also met the "expedient" criteria with ease, requiring under half a minute for data acquisition and analysis, not including mount and dismount times of the real device over the probe.

Scope and Limitations

Although the example above was for an environment that could accurately be described as piecewise continuous, the same algorithm could be used to parameterize nonlinear characteristics and features. For example, to characterize backlash a vibratory trajectory might be applied, and force hysteresis measured. Stiction could be measured by repeatedly stopping, then starting and measuring breakaway force at different positions and breaking rates.

[2] See the exeriment methods described in MacLean, 1996, for a discussion of the problem of measuring emulation fidelity, and experiment models for analyzing an emulation's success.

The operative term, however, is "parameterize". This algorithm measures parameters for an assumed model; it does not create the model structure. It can identify zero parameter values, indicating that the parameter was not a significant model component, but it cannot identify a parameter of which it was not informed. If the model provided does not describe the system well, this algorithm may misattribute impedance components (e.g., include all leftover impedance in its estimate of $\hat{M}(x)$ in the example above) and at best provide a statistic indicating the poor quality of the model fit.

Further limitations were inherent in the physical hardware which was used. In this realization, it was necessary to mount the real device over a non-portable probe, and the probe itself could only push. This limited the types of features it could measure; hysteresis would have been problematic. There are many types of real environment ports for which a probe could be devised to both push and pull with little hysteresis or chatter, but making a probe to meet that specification for arbitrary environment ports would be a challenge.

This algorithm, along with other probe-type characterizations, may be extended to parameterize some models such that interaction with the virtual mechanism could take place at other points than that probed during identification. This will require that the model structure resemble the real mechanism structurally, and that the excitation of the mechanism be rich enough to capture the entire mechanism's kinematic and dynamic complexity.

In future development, other more general methods of system identification may be brought to bear on this problem. For example, a frequency-rich stochastic input might be applied instead of the successive narrow-band trajectories, with less constraining assumptions about structure; it would require a push-pull probe. Stochastic inputs will be less useful in environments which are geometrically small (e.g. a few millimeters) with hard end stops, and/or contain narrow discontinuous zones of local continuity — with the discontinuities being the features of greatest interest.

Ideal Conditions for Use

This method is most useful for systems which are difficult to parameterize by simpler means, and for which a reasonable estimate of model structure can be made. It excels at geometric location of subtle and higher order changes in parameters which can be felt but are difficult to decompose perceptually. It simplifies the segmenting of a model into regions by collecting and analyzing a clear multi-layered picture of system impedance. It is less useful when there is a straightforward alternative or when the model structure is highly uncertain. Because error is cumulative, accuracy drops as more parameters are attempted; even in the above linear example, the mass evaluation will be less trustworthy than the stiffness and damping estimates. Mass characterization is thus a good example of when not to use this technique — mass is both a high order parameter and usually can be conveniently and accurately weighed.

The algorithm may be strengthened in some cases by the addition of other sources of parameter information. In the mass-spring-damper example above, the device might be disassembled and weighed prior to the more general characterization. This value for $\hat{M}(x)$, more likely to be constant over the whole device range of motion than the other parameters, can then be utilized as a model input rather than output, and the second order characterization (trajectories of constant acceleration) used as a redundant estimate of the lower order parameters, with its information weighted if desired according to the emulation designer's confidence in the method.

APPLICATIONS

This algorithm and its variants should prove useful in future efforts in creating haptic virtual environments, and in other applications as well.

Emulation of Real Environments: As this field grows and matures, the applications where correspondence between virtual and real environments is required will expand; and in many of those cases, an accurate and objective means of identifying the real environment will be required. An example is the hands-on surgeon training system whereby standard procedures are taught safely and haptically to novices, and unusual specialized techniques demonstrated to expert surgeons. Another is NASA's need to train astronauts to handle tools in unfamiliar environments. The tool can be characterized, the strange environment modeled, and the astronaut's dexterity trained and evaluated virtually.

Any Haptic Virtual Environment: Even when no direct correspondence to a real environment is required, learning how to produce difficult real haptic features will lend insight into creating richer virtual worlds in general.

Production Quality Control: Haptic characterization may find use as a quality control tool in a manufacturing environment; for example, to ensure that all the dashboards coming off the assembly line feel as intended. Imagine the QC inspector carrying around a handheld device and spot-checking dashboards by holding it up to individual switches: a rating comes up on the small LCD display and is stored in memory while a beeper sounds if the interface is out of specification.

Psychophysics Research: A mechanical identification procedure may prove valuable in improving our understanding of haptic perception, by providing an accurate and independent measure of environments at the physical point that people feel them.

[1] K. Dandekar. *Role of Mechanics in Tactile Sensing of Shape*. PhD thesis, Massachusetts Institute of Technology, Cambridge, MA, June 1995.

[2] C.R. Dohrmann, H.R. Busby, and D.M. Trujillo. Smoothing noisy data using dynamic programming and generalized cross-validation. *Journal of Biomechanical Engineering*, 110:37–41, 1988.

[3] J.K. Gotow, M.B. Friedman, and M.L. Nagurka. Controlled impedance test apparatus for studying human interpretation of kinesthetic feedback. In *Proceedings of the 1989 American Control Conference*, pages 332–337, Pittsburgh, PA, 1989.

[4] L.A. Jones and I.W. Hunter. Influence of the mechanical properties of a manipulandum on human operator dynamics. I. elastic stiffness. *Biological Cybernetics*, 62:299–307, 1990.

[5] D.A. Kontarinis and R.D. Howe. Tactile display of vibratory information in teleoperation and virtual environments. *Presence*, 1995.

[6] S.J. Lederman and R.L. Klatzky. Extracting object properties through haptic exploration. *Acta Psycholigica*, 84:29–40, 1993.

[7] S. Liu and H. Asada. Transferring manipulative skills to robots: Representation and acquisition of tool manipulative skills using a process dynamic model. *ASME Journal of Dynamic Systems, Measurement, and Control*, 114:220–228, 1992.

[8] K.E. MacLean. *Emulation of Haptic Feedback For Manual Interfaces*. PhD thesis, Massachusetts Institute of Technology, February 1996.

[9] K.E. MacLean and W.K. Durfee. An apparatus to study the emulation of haptic feedback. In *Dynamic Systems and Control Division*, volume 57-2, pages 615–621. Proceedings of the ASME International Mechanical Engineering Congress and Exposition, San Francisco, CA, 1995.

[10] T. Milner and D.W. Franklin. Two-dimensional endpoint stiffness of human fingers for flexor and extensor loads. In

Dynamic Systems and Control Division, volume 57-2, pages 649–656. Proceedings of the ASME International Mechanical Engineering Congress and Exposition, San Francisco, CA, 1995.

[11] J. Morrell and J.K. Salisbury. In pursuit of dynamic range: using parallel coupled actuators to overcome hardware limitations. In *Preprints of the Fourth International Symposium on Experimental Robotics, ISER'95*, Stanford, CA, June 1995.

[12] X.D. Pang, H.Z. Tan, and N.I. Durlach. Manual discrimination of force using active finger motion. *Perception and Psychophysics*, 49(6):531–540, 1991.

[13] G.A. Pratt and M.M. Williamson. Series elastic actuators. Artifical Intelligence Laboratory, Massachusetts Institute of Technology: unpublished, 1995.