To appear in Symposium on Haptic Feedback in the Proceedings of IEEE Robotics and Automation (ICRA'2000), San Francisco, CA, April 22-28, 2000.

Designing with Haptic Feedback

Karon E. MacLean

Interval Research Corporation
1801 Page Mill Road
Palo Alto, CA 94304 USA
karon@kmaclean.com

Abstract

Haptic feedback is a design element for human-computer interfaces, and this paper discusses when and how it can be used to best effect in interactive applications. It begins with consideration of the unique attributes of the touch sense in physiological and psychological terms, and the nature of information and control that touching provides. It reviews where active touching helps, by setting forth the forms it may take and important parameters that describe it; and evaluates the specific benefits it offers to contemporary interface problems. It ends with a proposal for a simple interaction model that emphasizes holistic design principles, and highlights issues that arise in the process of creating specific haptic interfaces.

1 Introduction

Haptic feedback is one of many available interaction mediums, with special properties that suit it uniquely to some contexts. Here we explore what touching is good for, how it can help in contemporary interaction applications, and offer a holistic view of how to use it.

Concerned with a creating a successful interaction rather than to using haptic feedback, an application designer will take a "top-down" approach. This begins with a need – to provide an effective interface to a given application – and finds a solution from a suite of technologies and methods. However, this implies knowledge of those tools and their affordances, and this is where we hope to help. Section 2 examines different aspects of physical interaction in the real world; Section 3 considers where and how active feedback can enhance interactive tasks. Section 4 concludes with a discussion of attributes that good haptic design should exhibit.

2 Understanding Physical Interaction

Touch has unique attributes best understood by reviewing how people use it in the natural world, considering psychophysical, cognitive and affective perspectives.

Special Qualities of Touch

Bidirectionality

The haptic sense is physically and neurally co-located and coordinated with motor functions, and the term "touch" commonly encompasses intention, manipulation and gesture as well as perception. Much of haptic perception relies on active exploration (Lederman & Klatzky [5]), which in turn is a form of gesture; and these integrated pathways allow fast reflexive motor responses to haptic stimuli.

Social Loading

Touch is intentional, socially invasive and committing. By reaching out to touch, we reveal our intentions, enter others' personal space and violate taboos. We expose ourselves to physical danger as well as pleasure and information. Because of its intimacy, social touching is salient and immediate.

Gesture and Expression

Some believe that verbal language evolved first from physical gesture, and is thus related to touch (e.g. Corballis & Lea [2]). Through touching, we convey functional signals (e.g., with a peremptory rap on a table) as well as emotion (by pensively stroking a familiar object, or clutching it tightly).

Multi-Parametered

As with vision and audition, touch has many qualitatively distinct components; this information is integrated with input from other senses to form a complex impression (Katz [4], Rock [8]). Parameters include force and pressure, moisture, temperature and spatial and temporal textures; each one can be further subdivided. Texture can be hard, rough, sticky, and wet. Even this is insufficient: the list does not yet capture, for example, the complex three-dimensionality of fur. The qualitative variety in haptic sensation makes it particularly challenging to classify and reproduce.

Resolution and Associability

Touch affords precise control and discrimination, but is vague compared to other senses in facilitating recall and association of absolute and relative resolutions. The finest scratch on a glass surface triggers a tactile reaction, and we can discern subtly different grades of sandpaper. But it is much more difficult to memorize and name those sandpaper grades, compared with color hues.

Reasons for Touching

We initiate or sustain a touch with many intentions. People are generally more cautious about what they touch than what they will look at; the designer must remember that they may have a choice. A large part of the task is thus anticipating, directing and accommodating a potential user's preconception of what the interaction will do, and what the experience will be like.

Motivations

We touch because we intend to *do* a task, *probe* an object for its state or qualities, *communicate* a message, *poke* something to elicit a reaction or *verify* that an action is completed. In more recreational circumstances, we expect to *enjoy* aesthetic pleasure or comfort, *fidget* to relieve tension, or *connect* physically or emotionally with another person or other living thing.

Inhibitions

Most often, we avoid a particular touch through a perception that it would be dirty, painful, forbidden or too intimate. Beyond this, many people (often culturally associated) are "haptically challenged": they do not generally find touching natural, informative or pleasant.

Information Available from Touching

Touch is the principal contributor to a number of high level, integrated perceptual functions:

Assessments of an object's dynamic and material properties. The particular information we seek (for example, texture versus weight) influences how we approach and handle the object [5].

Verification of engagement and completion. This is available as the satisfying "ka-chunk" from a button snap and an automobile shifter slipping into gear.

Continuous monitoring of ongoing activity and gradual doneness. The surging rattle of a vacuum cleaner sucking dirt, a pepper grinder's crunch and a pencil sharpener scraping wood relate progress and completion.

Building mental models for invisible parts of a system. We form hypotheses of its function, and physically probe to test them. A preconceived model influences both the final impression and the manner in which we explore.

Judgements of other people. A handshake is both a social gesture and a test or affirmation of a social hierarchy.

3 When Active Touching Helps

Where we can expect haptic feedback to be valuable? We should consider both the unique affordances of the haptic sense (push), and situations where other sensory channels are overloaded or otherwise unsuitable (pull). Haptic feedback is often most effective when associated with other sensory modalities and must be designed in conjunction with them.

Mediums of Tangibility

Tools and Textures

Physical objects usually exhibit a one-to-one correspondence of form to feel. Although one handle may be employed in multiple ways, most hand tools of yore were customized for a job. The result, illustrated by the contents of craftsperson's toolchest, is a wide range of shapes and many variations on a theme. Individuals often modify a tool to fit their own needs. Some tools are rarely used, but occasionally essential.

Textures on physical objects and surfaces serve many functions, some dependent on their non- or slowly changing nature. They provide friction for grip, slickness for motion, aid recognition through their distinctiveness, and indicate wear. Textures may be created deliberately or as artifacts of production or use; they may be informative or designed to enhance dexterity.

Haptic Language

Most people understand haptic language – the lexicon and syntax of affective communication through touch – intuitively and effortlessly, absorbing its grammar in youth when they learn other languages. Touch shares many attributes with visual gesture, in conversation and dance, but it has not been studied or linguistically codified to the same extent – a perquisite for machine recognition and synthesis. The first step is to create a lexicon of distinguishable haptic symbols in the form of an orthogonal haptic parameter space. Auditory research offers inspiration, e.g. the use of multi-dimensional scaling to find independent axes for timbre (Wessel [11]).

One can also speculate on the existence or learnability of a direct lexical relation to other sensory mediums, a general kind of synesthesia (a rare hot-wiring between senses wherein a stimulus to one elicits a percept in another - Cytowic [3]). The visual excitement generated by a hot color like red, for example, might translate to an abrupt, racing, hot haptic sensation. Again, this work could borrow from auditory research; "earcons" are abstract, learned auditory linguistic elements (Blattner et al [1]), as opposed to auditory icons whose connotations are experientially derived (e.g. the sound of a telephone ringing). Successfully matched multimodal sensations can be more salient than unimodal stimuli; and by understanding the translation, stimuli in different senses can be substituted.

Synthesized Haptic Feedback

"Haptic feedback" has come to imply computer control over the tactile or kinesthetic properties of a physical interface, permitting realtime representation of a changing virtual or remote environment rather than a specific, constant handle. However, power-supplying actuators are just one means of modulating feel. A computer-controlled brake's passive dissipation of a user's input energy has advantages of stability and potentially lower power consumption. Even more exotic is the "parasitic" haptic display, which absorbs and stores a user's own energy and offers it back in the form of active haptic feedback at a later time. A trivial example is a spring-loaded button that is depressed and locked, then released at a later time either automatically or triggered by another user action.

Attention has focussed on kinesthetic and vibrotactile haptics; but these afford a small subset of the haptic sensations available in the real world. By *multihaptics* we refer to the seamless and spatially overlaid integration of different haptic modalities, including temperature and moisture and greater variety of texture and shape.

Mediating Haptic Interfaces

To date, most haptic interface research has been devoted to directly exploring or manipulating static or dynamic virtual or remote environments; information is transmitted directly between user and environment. A haptic interface may also *mediate* between users or between a person and a machine, in both abstract and affective ways (Snibbe et al [6]). These include extending haptic control and exploration to media that does not have a direct physical analog (e.g. a video stream or database), introducing affect to electronic applications and providing more personal connections between people separated in space, as well as stability in the face of time delays.

Potential Uses

Reconfigurability

Whereas a manual interface without actuation or computer control (e.g. a computer mouse) can also provide benefits of physicality and continuous control, both actively and passively actuated haptic interfaces can change their feedback in response to the environment they display and control. This might be as simple as the ability to alter the number of detents around a haptic knob to reflect different densities in the controlled media; or as sophisticated as a 6-df robot used to interact with a complex dynamic virtual environment.

Handles for Continuous Control and Monitoring

Control handles provide continuous, analog user guidance or intervention. Haptic feedback can reduce motor or visual strain when the manipulation is exacting or prolonged. It can offer selective, suggestive guidance with a cue that the user can smoothly and variably over-ride, as well as a gentle resistance against which it is easier to control motion.

Expressive control of a variable or process usually employs continuous input, for example in sketching a visual image or musical melody. Haptic feedback can further enhance this by mediating the input with a dynamic interaction model, increasing the variations possible within a given medium and providing an avenue for stylistic experimentation [9].

Low-resolution, low-attention monitoring is a good candidate for haptic feedback. Changes in a haptic landscape – features and discontinuities – are more salient than absolute values in both temporal and spatial domains and don't require memorization or recognition.

Teaching, training and guiding of manual tasks and gestures can be facilitated by "intelligent" haptic systems able to diminish the teaching cue as a student learns.

Buttons for Discrete Control and Information

Differentiation and identification of discrete objects, surfaces and boundaries is aided by recognizable, associable tactile properties, either static or dynamic. These button-like objects may be static physical artifacts whose passive haptic properties have been carefully designed, or they might possess dynamic haptic behaviors to help relieve the semantic loading placed on visually indistinguishable tagged objects.

Imposing discretization on continuous input can relieve the strain imposed by generic I/O. An active haptic mouse allows a user to feel the edges of windows and pull-down menus without falling off.

A user can be *notified* of events unobtrusively and with an informative range of values. For example, wireless devices with haptic capability impart an incoming call or an alarm without bothering neighbors.

A device's *failure* or an action's *confirmation* is often communicated via passive haptic interaction in real mechanisms. These subtle cues can be incorporated into sophisticated electronic interaction.

Touch is a locus for *reflex-rate user reactions*, measured in milliseconds. Certain categories of manual tool control share this need, and haptic feedback could be used to elicit and transmit the user reaction from computer-supplied stimuli at these rates – e.g. surgical and musical instruments.

Affect and Communication

Haptic feedback can add *social context* to a socially sensitive or impoverished situation, e.g. computer-mediated connections between people or between people and computers in professional, personal and entertainment domains. Communicating affect or personal presence in a variety of forms may enrich such situations, and augment the sense of a shared experience. As we learn its language and build on its strong social and personal-space connotations, haptic and gestural feedback could be a central means to this.

Comfort and Aesthetics

Much of our natural touching gratifies urges purely of aesthetics and comfort – for example, stroking attractive surfaces and fabrics, and fidgeting with articulated objects. As mechanisms and natural materials give way to digital circuits and plastic, opportunities for such indulgence become scarce. Gratuitous addition of nice-feeling haptic qualities can immeasurable enhance the pleasure of interaction. What haptics can add includes:

Pleasant tactility that triggers a desire to touch, and then hold or stroke an object.

Satisfying motion and dynamics are best when they are also informative.

Ergonomics should avoid cramping, clutching, nerve pressure and heavy use of weak fingers.

Bidirectional coupling with environment provides a sense of being inside, rather than directing it from without.

Muscle memory can be used to structure frequent, patterned tasks into stylized or abbreviated gestures.

Personalization makes a device special and suggests ownership and value.

Dealing with Complexity

The volume of information and things and devices people routinely encounter can overwhelm them. There are many sources of complexity – too many, hard to distinguish or remember, easy to lose, takes too many steps, can't tell how it works. They are exacerbated by electronic technology, which tends to have many operations, propagate ranks of buttons, hide functionality, rely on menus and screens, and to be wireless and buried in the couch. They rarely help the user in forming a useful mental model of the system being controlled.

A suite of well-designed embedded physical interfaces, including active and passive haptic displays and tagged tangible objects (Ullmer [10]), can help with different subsets of these problems. Context-sensitive active haptic feedback in a knob on a handheld controller permits a single analog input to be clearly redirected, if the result of the control action is consistent in the environment and in the haptic feedback. Feeling a virtual representation or "map" of an electronic system's operational model can help a user understand how it works. Haptic feedback can offer clues as to what a user's options are, through constraints and gentle guidance. Tactility can be used to differentiate sets of buttons. Sequences of discrete steps can be merged into a single fluid continuous control gesture. Active objects can transport electronic tool use away from the desktop computer.

Other Areas of Value

The techno-literati are not wired; they are wireless and portable**.** *Wearable controls and information displays* are active research areas, and good manual controllers will play a key role in their ultimate usability. Currently haptic feedback in wearable devices is challenging because of size and power requirements of most conventional haptic actuation techniques, but with creativity and constraints there are ways around this.

Biomedical and prosthetic applications include augmentation, filtering and otherwise supporting manual activities by the variously disabled. With the growing prevalence of keyboard-and mouse induced repetitive strain injury, this research area may hit mainstream.

4 Designing it In

A Model for Multisensory Interaction

The foregoing organizes experiential evidence of the ways touch connects us to the physical world, and suggests how touch could be utilized in human-machine communication. Here we outline an interaction model that includes all sensory modalities.

Figure 1 is a generalized view of a sensory interaction as a multi-layered and multi-modal structure. The user is on the outside, the manipulated environment at the core, and layers of physical interface and interaction models in between, wrapped onion-like around the environment. The notable features of this representation are:

- Acknowledgement of the integral multi-sensory aspect of most haptic interactions.
- The explicit presence of an interaction model between the physical hardware and the environment being manipulated or perceived.

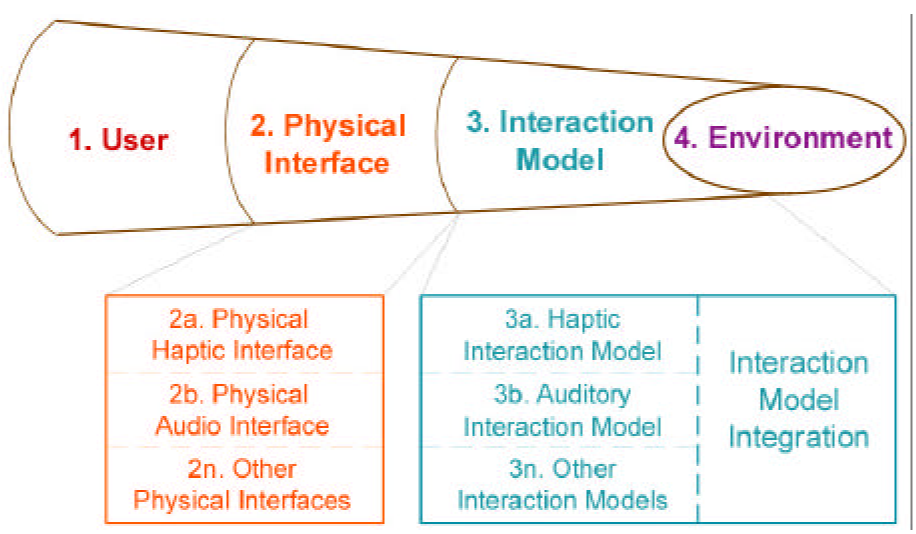

Figure 1. A model of multi-sensory interaction: the environment (Layer 4) is wrapped with an interaction model (Layer 3), then a multisensory interface (Layer 2).

The latter allows an arbitrary relation between user and environment, whether direct or abstract. Whereas prior emphasis has dwelt on creating the outer hardware layer, there is much to be done by creatively disposing the interaction model into abstract forms.

The User (Outermost Layer)

What the user is requires little definition. How the user is disposed, or would like to be, is critical to an interface's design. This includes issues such as the level of attention he wishes to devote to controlling this environment, what senses and limbs are available for the task, and whether he is likely to be disabled in any way or to have specially trained abilities. Most importantly, what are all the things he wants to do? Standard task analysis methods should be applied for specific contexts.

The Environment (Innermost Layer)

The environment is whatever the interface is intended to observe or manipulate; for example, a CAD representation of a physical mechanism or a solid-body model. It could be the lighting and temperature of rooms throughout a house, your account balance at an ATM machine, a video or audio stream, a database, or a Windows screen. The environment is distinct from the interaction model and its components and may be represented in arbitrary ways, unrelated to its own form. If the environment is a 3D model of a crane, the literal interaction model of directly touching and manipulating its moving parts is just one of many possible.

The Physical Interface (Layer 2)

The physical interface consists of the electromechanical transducers and displays that accept input from and provide output to the user for all the sensory modalities. For a graphical interface, this may be a CRT; for an auditory interface, a speaker. The haptic physical interface is the I/O for the haptic display and motor control.

The Interaction Model (Layer 3)

The interaction model defines the relation between user and environment, in ways described below. Each sensory mode has a submodel that generates its display. Likewise, each input modality (touch, voice, etc.) is processed in the submodel, producing signals that are integrated and transmitted to the environment. Each user-input and output modality may fill a reassignable role relative to the integrated interaction model.

Abstraction in User-Environment Mediation

Haptic mediation between a user and an environment can be pitched at different levels of abstraction.

Direct Manipulation

At the most literal, an environment can be rendered, probed or palpated through an interaction model that mimics the environment itself. Such a haptic interaction allows the user to directly "feel" the environment or signal, with minimal interposed abstraction.

Container Manipulation

A simple, relatively automatic hierarchical layer of abstraction is interposed between environment and user. The user feels containers or bins of the raw signal or model; for example, bumps corresponding to successive frames of a video signal, as opposed to the raw bits and bytes coming down the line.

Annotation

Editorial content is added to a signal or model environment in the form of annotations, which a user can perceive haptically. The annotations are generally linked to discrete locations or components of the environment, rather than becoming a general property of the interaction model.

Mediating Dynamic System

The sensory-mode interaction models and the block that integrates and coordinates them comprise an arbitrary, abstract dynamic system used to manipulate and observe the environment. This interaction model might bear no direct physical relation to the environment. For example, the interaction model for a video environment might be a spinning virtual mass. The user interacts with the video by spinning up and braking the mass, whose rotation is linked to the visual frame rate.

This last and most abstract interaction model offers powerful, intuitive control of many environments that might be difficult to represent literally. That its structure is not tied to any physics of the environment empowers the designer, but may also entail greater challenge to make the interaction intuitive.

As with other mediums of interaction design, it is common to employ metaphor to find useful abstractions and generate intuitive tangibility in computer-mediated processes. For haptic manipulation, many of the relevant metaphors will relate to conventional manual tools, actions, materiality and objects.

Other Challenges for Haptic Design

Discrete and Continuous Control

Haptic feedback can be useful for both discrete and continuous regimes of manual control, but it will generally be most valuable when the latter is required. When a task has components of both (e.g., the need to discretely change mode or content of an environment as well as continuously manipulate that content), designing the affordance for both and the transition between them to be intuitive and seamless can be nontrivial (MacLean et al [7]).

This tension can also be an opportunity for a revised interaction model. For example, in the case of browsing streaming media such as video, the environment's gradual transition from a discrete (individual frames) to a continuous regime based on frame rate could be related to the "freezing" of granular elements into a rigid body. In the frozen regime, the rigid body can be shoved, spun up like a flywheel, and perhaps stretched as a coherent elastic body. Alternatively, the discrete phase might be seen as frozen, then "melting" into a fluidic continuous phase; fluid metaphors such as spraying and pouring would then come to mind.

Displaying Interaction Potential

A requirement of any good interaction model is that it makes clear to a user not only *how* to switch an activity or elicit a behavior supported by the model, but that the potential of doing this exists and what will happen as a result. A haptic interface might render this more challenging than is usual, because it lacks the textual "fixing" available with menus and toolbars and visual icons. Further, it has yet to acquire a history of past interfaces and expectations; any user knows that "ctrl-V" will paste the clipboard into a Windows application.

The strong use of metaphor and suggestibility as well as simplicity is a good way to start, since it invites using the interface in the same way as the implied physical object. As custom and familiarity with haptic interfaces grow, it is reasonable to expect that language and conventions will develop to make this problem easier. Of course, the emergent conventions might be the wrong ones – beware another QWERTY keyboard.

Embedding Haptic Interfaces

Physical design consistent with an abstract interaction model leverages intuitive haptic interaction. For example, a spinning-mass interaction model for video browsing is natural when the video is displayed with low latency, and the haptic display looks and tactually feels like a wheel. If the haptic display looks like a pen-probe or a mouse, it may be harder to figure out.

This diverges from the more prevalent approach to building haptic displays, which for excellent reason seeks to provide general-purpose access to graphic displays. However, these interfaces are often expensive as well as highly generalized. When the requirement of generality is relaxed, a cheaper, simpler special-purpose haptic display can be created with a handle crafted for a given task. We hope to see this approach spread on the coattails of embedded controller technology.

Tight Sensory Coupling for Perceived Control

Some of the best-feeling synthesized haptics derive not from the specific qualities of the haptic feedback, but from the low latency by which it is linked to environment manipulation and other sensory displays. We call these "tight-coupled" displays, and value the sense of presence and control they offer the user. Tight coupling can be achieved in various ways, through the use of realtime software architectures, high interprocess communication rates and code customized at a low level. These methods need not be expensive (MacLean et al [6]).

Some applications particularly justify the maintenance of low latency communication among interaction model elements; for instance, musical and drawing controllers, which rely on expressivity and crisp response. However, the effect is so satisfying that we hope it will become a minimal performance metric for multi-sensory displays.

Acknowledgements

This paper summarizes insights developed at Interval Research Corporation by the haptics team and its friends and guests. In particular Scott Snibbe, Bill Verplank, Rob Shaw, Jayne Roderick, Perry Cook and Durrell Bishop have helped inspire and instantiate them.

References

[1] M. M. Blattner, D. A. Sumikawa, and R. M. Greenberg, "Earcons and Icons: Their Structure and Common Design Principles," Human-Computer Interaction, vol. 4:1, pp. 11-44, 1989.

[2] M. C. Corballis and S. E. G. Lea, "The Descent of Mind: Psychological Perspectives on Hominid Evolution," : Oxford University Press, 1999.

[3] R. E. Cytowic, The Man Who Tasted Shapes: A Bizarre Medical Mystery Offers Revolutionary Insights into Reasoning, Emotion, and Consciousness. New York: Putnam, 1993.

[4] D. Katz, The World of Touch. New Jersey: Erlbaum, 1925 / 89.

[5] S. J. Lederman and R. L. Klatzky, "Hand movements: a window into haptic object recognition," Cognitive Psychology, vol. 19:3, pp. 342-368, 1987.

[6] K. E. MacLean and S. S. Snibbe, "An Architecture for Haptic Control of Media," in Proc. of the 8th Ann. Symp. on Haptic Interfaces for Virtual Environment and Teleoperator Systems, ASME / IMECE, Nashville, TN, DSC-5B-3, 1999.

[7] K. E. MacLean, S. S. Snibbe, and G. Levin, "Tagged Handles: Merging Discrete and Continuous Control," in Proc. of the Conference on Human Factors in Computing Systems (CHI '2000), The Hague, Netherlands, 2000.

[8] I. Rock and C. S. Harris, "Vision and Touch," Scientific American, vol. 216 (5), pp. 96-104, 1967.

[9] S. S. Snibbe, K. E. MacLean, J. B. Roderick, R. Shaw, M. Scheeff, and W. Verplank, "Haptic Metaphors for Digital Media," Interval Research Corp., Palo Alto, TR IRC #1999-071, 2000.

[10] B. Ullmer, H. Ishii, and D. Glass, "mediaBlocks: Physical Containers, Transports, and Controls for Online Media," in Proc. of Siggraph '98, 1998.

[11] D. L. Wessel, "Timbre Space as a Musical Control Structure," Computer Music Journal, vol. 3:2, 1979.