TUTORIAL

Do It Yourself Haptics: Part II

Interaction Design

BY KARON E. MACLEAN AND VINCENT HAYWARD

This article is the second of a two-part series intended to be an introduction to haptic interfaces, their construction, and application design. Haptic interactions employ mechanical, programmed physical devices that can be used for human-computer communication via the sense of touch. In Part I of this series, we focused on the devices themselves: the classes of hardware schemes currently available or envisioned, the software components that drive them, and specific examples that can be built on the kitchen table. Here in Part II, we broach a topic that is coming into its own; between the vision of a particular utility that haptic feedback theoretically should enable and the hardware capable of delivering the required sensations is the problem of designing the interaction in a usable way.

Introduction

Haptic technology has hit the mainstream. In 2000, there weren't that many people who knew that the word *haptic* definitely did not refer to a liver dysfunction. By 2004, any self-respecting gamer had it in a joystick at home, and cell phones buzzed. Today, these devices already show the potential to transform many specialized tasks, and the vision of embedded, haptically enabled devices soon dominating our everyday existence is shared by a guru of human–computer interaction (HCI) [71]. It is an inevitable development, despite considerable technological challenges. Our information age has taken the path of networking and abstractions; yet, evolutionarily, we are physical animals dependent on touch to function and communicate. As information technology matures and continually becomes more complex and intrusive, its intangibility and remoteness (action at a distance) become more obvious flaws. Haptic technology offers a solution—if we do it right.

The haptic sense, comprising taction (mediated by the skin) and proprioception (our conscious or unconscious experience of body movements and forces), is often observed to be special in its close association with motor channels—one perceives and acts in tight integration. Today, it has another imputed virtue: that of not simply not being either vision or audition. Contemporary computational interfaces have saturated our eyes and ears. There's not much communication bandwidth left there, whether one is an automobile driver, an urban pedestrian, or a medical professional in the operating room. It is therefore common to suggest that beyond its role in providing tangibility and real-world fidelity, the touch sense is another potential information conduit. Thus, we see at least two distinct and major role types for haptics in

- restoring tangibility to digital interactions, with functional and aesthetic potential
- offering an additional communication conduit, providing we recognize the importance of attentional design and the overall user environment and its loading.

We'll be going into these aspects, which have many facets and can overlap, in more detail.

Why Interaction Design Matters

There are not many computer users today without a collection of stories of user interfaces (UIs)—generally graphical, as that is what we are surrounded with—that have annoyed, confused, or stymied them. The frequency of these incidents has unfortunately not diminished with time and experience nor are they, in most cases, due to limits in the extraordinary graphical display and back-end hardware available today. They

Digital Object Identifier 10.1109/M-RA.2007.914919

1070-9932/08/$25.00©2008 IEEE

are, rather, the intersection of bad UI design by untrained and unsupported application creators and paying customers who clamor (or respond to marketing) for features and style rather than recognizing and valuing usability. These problems are exacerbated by the remarkable number of technologically supported tasks that we now tend to do at the same time. It is like being treated for multiple ailments by several specialist doctors who cannot or will not coordinate with one another—leaving the patient/user to sort out the impossible conflicts alone.

As some forms of haptic technology depart research labs as commodities, it is exhibiting a similar phenomenon. It is becoming technically feasible to integrate haptic feedback into everyday devices, but it is also easy to misuse it—far easier, in fact, than to use it well. Good UI design is hard. It's not just a need for formal training and experience, which helps, but much of what is taught is really just a codification of common sense. The tough part is taking the time, space, and money in a given design cycle to

1) truly understand the user's experience, problems and needs—the whole context of the interaction; this happens by observing and talking to said users
2) base design prototypes (ideally, a few very different approaches) on thorough knowledge of relevant human capabilities, in terms of perceptual, cognitive, and motor attributes. These, again must be related to the entire context; if a user is doing many things at once, that means their resources are not fully available for your task
3) verify and iterate on a design prototype through user testing, rather than relying on a designer's guess of what will work
4) allow the UI design to influence the rest of the system's design to support an optimization of the user's experience (as opposed to, say, a feature list created by the marketing group, which is longer than the list of a competitor's product). Sometimes, a good UI design will indicate a change in a device's physical form factor. If the UI has been slapped on as a final step, this will probably be impossible.

These basic principles of good HCI design are all the more important when the modality is one that people are not accustomed to using in this way and, furthermore, one that is often being layered on top of whatever else the user is already seeing or hearing. It's a perfect storm for sensory and cognitive overload.

This article's primary goal is to provide some basic heuristics and examples for avoiding that storm, and instead offer a path for integrating haptic feedback into the mix of the user experience in a way that will help.

Overview

In the remainder of this article, we'll start by considering the mapping between the crosscutting roles, which haptic feedback is thought to serve, in many different kinds of application spaces and, conversely, the human abilities and limitations that must be recognized, targeted, or supported as these roles are developed (see the "Usable Roles for Haptic Feedback" section). We'll provide some design guidance, which is especially relevant to haptic interactions (see the "Haptic Interaction Design Practices" section), and then close with a pair of case studies that illustrate contrasting approaches to actually doing it (see the "Design Case Studies" section).

Usable Roles for Haptic Feedback

Previously, we listed some very broadly defined several potential *roles* for haptic feedback. On a closer look, here, we take a different cut. In each of the several categories (the list is certainly not exhaustive), we will consider haptic value in terms of functionality, emotion and aesthetics, in search of ways in which it can improve task performance or expand capabilities, allow us to communicate through technological conduits, or make an interaction more pleasurable and satisfying. Some of the categories relate to control, i.e., the closely coupled perception-motor action loop referred to earlier. Others are more sensory in nature, e.g., tactile messaging where the skin is used as a display surface, but the user's response might be less direct—e.g., a thought or a directed look. For additional background, we refer the reader to some recent comprehensive reviews of human sensory, cognitive, attentional, and motor abilities, which [63] summarizes in the context of interaction design.

Naturalistic Interactions

A common theme in the following discussion is to relate new potential functionality to natural, i.e., ecological, touch interactions in the nontechnological world. Our sensorimotor equipment and social wiring are likely to be well evolved or conditioned to handle the things we do naturally, comfortably, and with easy precision in this domain.

This is not an adage to follow slavishly, however. There are many examples of humans picking up new technological skills with apparent ease, despite a lack of obvious evolutionary preparation (driving a car, typing, and perhaps most remarkably, text messaging on tiny cell phone keyboards). We already see evidence of this here, e.g., in human acuity in abstract tactile message decoding, an unnatural act that will come back to haunt us with stress and damaged thumbs? Perhaps only time will tell.

Multimodality of Haptic Interactions

Haptic design is nearly always a multimodal design. The touch sense is generally used in conjunction with other sensory modalities, whether their roles are to reinforce the same task or to handle different tasks performed at the same time. Touch-derived input plays a unique role in this context, and theories continue to develop on how sensory information is integrated and how conflicting information is resolved. The emerging answer is that relevance of the source to the task matters along with the source's trustworthiness [30].

Precise Control: Force Versus Position

We will start with a low-level attribute of coupled perception-action applications (usually involving force feedback), because of its far-ranging and often overlooked consequences. The sensation and control of absolute position is easily perturbed—try to reach out and touch a specific point in space with your hand while turning your gaze away and without groping for landmarks. Conversely, we're quite skilled at detecting and producing small variations in force resistance. This is seen in a comparison of natural, ungrounded human gestures (conversational emphasis, demonstrating emotion, or indicating a relatively discrete-valued command—stop, come, look over there) with those that entail resistance (almost any kind of tool use, from chopping vegetables to writing and painting, maintaining a desired pressure on an automobile throttle, precisely controlling a violin string's vibration). For humans, precision requires resistance.

The implication for design is that grounded resistance—something solid to push against—is desirable for most kinds of precise tasks. It is imperative to remember this when choosing what will be displayed, and the tasks best suited to haptic augmentation. To implement this principle, resistance could be provided by a programmed force feedback system or, alternatively, by a passive ground (e.g., a tabletop) with nongrounded feedback (such as impulses or vibrations) supplying the programmed feedback. In this latter case, in pushing against a stiff surface, the user's input will be isometric (without measurable motion), and so position sensing cannot be used to measure user intent. However, pressure might be more suitable.

When precision is not needed, and broad expansive gestures are appropriate, then nongrounded systems (such as a limb-mounted tactile display) might be more appropriate.

High-Fidelity Rendering and Model Creation

The role for haptic feedback, which has received the greatest research attention to date, is the creation and literal haptic rendering of what we see on a graphical display. These efforts have been dominated by surgical simulation and remote surgical procedures. Because of their substantial coverage elsewhere ([14], [22], [47], [56], [87]; see also Part I of this tutorial), we will not discuss them in detail here but place them in context with other uses and relate this role to human attributes.

A dominant and fairly unique aspect of these applications is their need for high fidelity to real-world analogs, so as to recreate a specific task environment—e.g., for training, or for actually conducting a remote or virtualized version of a task, which was once performed physically. Because of this direct tie, high fidelity rendering obviously borrows heavily from haptic interactions in the real world. In some cases, the real world case can be improved upon (for example, a tool geometry that is awkward or misscaled in reality can be reconfigured or magnified).

Obtaining satisfactory fidelity is one challenge, as discussed in Part I. The turing test of haptic rendering would be a user's inability to distinguish it from the real thing. In fact, this is currently possible for only a small subset of possible rendering targets, usually the more squishy ones, and thus usability can mean identifying and exploiting the limitations of the perceptual system to reduce the negative impact of system constraints. Another design direction is in augmentation, e.g., reconfiguring an operation or layering information atop a rendering such as signals or virtual fixtures (more about these are discussed later).

An additional element is the creation of the models themselves, which can be done through a variety of empirical and analytical, automated and manual approaches (a brief review is available in [63]). In particular, it is necessary to understand a user's perceptual attributes to specify the resolution, stiffness, and many other aspects of the model. In general, highly detailed and stiff renderings—exactly what you'd need to recreate many interesting physical systems—are difficult to stabilize, and the resulting artifacts destroy the illusion of realism [19]. Thus, the designer is often faced with a tradeoff between overall realism versus fidelity in shape detail, texture, hardness, dynamic response, and other rendering parameters. Alleviating this tradeoff drives much of the research in rendering techniques [56].

Finally, multimodal issues are almost always critical to attaining a realistic simulation result, in particular for renderings that need to convey high stiffness. In these cases, achieving visual-haptic and audio-haptic synchrony to perceptual limits will allow perceptual fusion of the information arriving on the different sensory modalities. Furthermore, the presence of the visual and auditory stimuli can significantly modify the user's interpretation of what they feel, allowing the use of less expensive or slower haptic hardware (e.g., [23], [44], [55], [105]).

Physical Guidance

Both force and tactile feedback can be used to provide direct spatial guidance to a user, either by leading with forces or orienting attention in a particular direction. Attentional orientation usually takes the form of applying a discrete signal to a body location, which then draws visual attention in the same direction, or providing an information-containing signal at a single location (which is discussed more in the following section). Guidance, on the other hand, implies a more continuous engagement that is usually delivered through grounded force feedback for motor skills or, with lower resolution, via distributed tactors on the body for applications such as vehicle steering. It can vary in precision and subtlety, for example, steering a car or aircraft, drawing a calligraphic character, or learning a surgical procedure. Force feedback guidance applications tend to vary across the spectrum of control sharing with the intelligent system (i.e., equally shared versus dominated by one or the other).

Training

In teaching applications, the user is expected to exactly follow the intelligent system's lead. The teacher or another human could be an expert system, and the latter is an instance of shared control or remote collaboration, which is also discussed more in the next section. These methods have been tested in applications ranging from calligraphic writing and surgical tasks to rehabilitation therapy for stroke patients. Haptic feedback has been shown to have value in the training of sensorimotor tasks, with improved performance in a real version of the task following inclusion of haptic feedback in a virtual-reality training segment [1], [69], when the real task has a force component. It has been further observed that visual training is better for teaching trajectory shape, but haptic guidance is more effective for temporal aspects [31].

There are many variants of implementing the construction of training forces. These include guiding the user along a predefined trajectory [2], displaying both the activating pressure and position of the teacher to the student (one indirectly) [51], and requiring the student to cancel a reversed target force [84]. More long-term learning strategies include monitoring the student's resistance and backing off as the need for guidance decreases. This also allows a simultaneous assessment capability [38], [53], [99]. These methods have not been directly compared with one another, and so at this point, it is difficult to evaluate their relative appropriateness in different situations. However, there seems little debate that the creation of motor programs requires realistic resistance to fully develop.

Shared Control

The notion of shared control refers to a cooperative balance of control between the user and the machine. An expert system has knowledge of the sensed and networked environment and the databases but does not know the user's goals. In this case, the system and user can jointly exert the forces that control the system. This concept is especially natural in steering contexts, where there is a single locus of control (e.g., a steering wheel or aircraft stick) that is intuitive to specify in a physical manner.

Telerobotics: Force sharing lies on a continuum of abstraction, which has at one end bilateral force-reflecting telerobots. These systems consist of a remote robot located in the work environment, connected through a network to a local robot of compatible kinematic design, which an operator moves, often wearing it as an exoskeleton and feeling forces sensed remotely and redisplayed locally. This scheme allows the local user to be sensitive to the impedance of the remote environment, with consequently improved dexterity and error management (an early instance is [48]; the beginnings of force sharing during teleoperation is illustrated in [98]).

Virtual Fixtures: The most common basis for shared control derives from the idea of a physical template for guiding a task by keeping it within specified constraints (e.g., a ruler for drawing a straight line). In a virtual environment, programmed forces provide the constraint [82]. Softening the guiding constraint turns this concept into mixed-initiative guidance: the user can choose to be guided or punched through to do something else. Many variants of control sharing using this concept have been tried ([34], [41], [52], [57], [73]; see [63] for a more thorough discussion). A sought-after metric is improvement in task performance while reducing visual demand, thus freeing attention for other tasks, and this has indeed been shown.

In extending these ideas to less predictable, real-world scenarios, however, there are additional complications. In particular, the reflexive dynamics introduced by the user can make them tricky to implement, e.g., oscillations can result from certain kinds of system disturbances [34]. Usable solutions depend on the task, but ideally they will build upon an as yet incomplete knowledge base deriving from both modeling of the user's reflexive and cognitive responses to control actions that are perceived as intrusive, and user testing in both abstract and reasonably realistic contexts.

Cognitive Factors: The user's mindset and awareness of the control balance is a variable to be managed. There are potentially negative side-effects, for example due to the operator's either over- or under-trusting the control suggestions or not understanding who is in charge at a given time [27], [40]. For this reason, it is crucial to manage the reliability of the expert system's signals. The idea of tuning the ratio of hits and misses for an expert system's detection and communication of crucial environmental events (e.g., dangerously close following of the car ahead [27]) and its effect on operator utility of those signals as well as overall efficiency has roots in multiple resource theory, recently updated in [25].

Remote Collaboration

When force communication is important, remote collaboration with another human in a physical task becomes a special case of shared force control (where the automatic controller potentially still plays an important role). This case is particularly interesting because, beyond the demonstrated need to feel the forces to perform a physical task, the existence of another human in the loop introduces social factors as well; and feeling ones' partner's forces appears to be an important parameter in facilitating this. It enhances the sense of presence and togetherness in the mutual effort [6], [85] and conveys the momentary degree of control balance between the partners [72]. In an explicitly social context, the nature of the force sharing impacts the sense of an interpersonal emotional connection [88].

Tactile Signaling in Multitasking Environments

Passive touch cues (which are presented to the observer's skin, rather than felt in response to active movements [36]) can be used for notification of events and to create relatively nonintrusive, ambient background awareness. Such cues can be delivered through a tactile display or overlaid on a force feedback signal being used for another function.

Typically, this kind of functionality targets *multitasking environments* where the user's primary attention, as well as visual resources and possibly hands, are engaged in another task (in fact, this benefit was foreseen very early on in the technology's development [79]). In this section, we'll therefore first mention issues relating to tactile design for multitasking, as well as typical methods and sites of delivery. We will then look at two major categories of tactile signals themselves: simple signals whose message comprises its on/off state (sometimes coordinated with its location), versus informative signals (haptic icons) that can vary in other parameters, e.g., amplitude or feel, and thereby encode additional meaning. Analogous auditory signals are a simple, consistent beep (perhaps directional) versus the diverse auditory icons we hear on modern computers whose specific sound means something—like an application opening, a device ejecting, or an e-mail arriving. Design in these cases is best based on some understanding of human multisensory attention. An overview, including references to other relevant recent work, can be found in [63].

Design for Multitasking Environments

To manage intrusiveness, tactile signals must be designed with variable salience: important events or urgent events/changes should register as louder than less important ones [16]. Furthermore, the user's interruptibility is not a constant, sensory adaptation aside. In the car, pulled-over versus engaged in a turn differ substantially in what kind of additional distractions the driver can safely deal with. In the office, some tasks require protection from routine interference, and yet certain events might always be important enough to come through. This entails two different needs, both active research areas.

Controlling Tactile Signal Salience: It is most desirable to control signal salience independently of potential content. In different contexts, a given event might be more or less important; and in some cases, context may be identifiable.

Parameters used to encode content may also vary inherently in salience. For example, in some schemes and for some display hardware, higher frequencies and/or amplitudes are perceived as louder than lower ones, yet these are the best parameters to vary to indicate different meanings—the change in output is easy to produce precisely and is clearly detectable by a human. Therefore, salience can be inadvertently confounded with meaning, with an unimportant signal more detectable and intrusive than a critical one. This incidence can be minimized with an up-front awareness of the stimulus salience and detectability patterns for a given display. While it is easy to determine relative salience (by itself) for a group of signals, e.g., using simple subjective ranking tests, due to this confound there is a need for design tools that efficiently aid this task *at the same time* as optimizing design of meaning.

Context Detection: The other part of the problem is detecting the user's momentary environment so that the appropriate salience can be used. The active field of sensor-based computing is devoted in part to detecting various aspects of the user context (e.g., location) [68], [76] and in modeling and detecting user mental/emotional state and interruptibility [32], [46].

Ambient Tactile Displays and the Human Body

Physical Configuration and Body Site: It is necessary for ambient tactile displays to be in continual contact with the stimulus site, so that signals will not be missed. Because the hands are often needed for more dexterous roles, the glabrous skin of the fingertips not always convenient as the delivery site, which leaves the less sensitive hairy skin [35]. Past examples, usually for simple signals, have used vests and belts [50], [95], [102], back [95], [106], and tongue [3], and relied on spatial encoding of meaning.

Applications and contexts where hands can be used for background display include the steering wheel [26], track point [15], mouse [16], [17], and increasingly, mobile devices [54], [60], [81].

Active and Passive Touch: More fundamentally, Gibson has argued that "passive touch ... is atypical of normal tactile perception and it leads the person to focus on the body surface" [36], whereas active touch is predominant in naturalistic environments where people are seeking information [86]. Considering that convenient ambient tactile delivery sites are generally less sensitive skin and that the information is intended to be nonattentive, it will be an experimental challenge to test the implication that passively received information display will be less effective.

Simple Tactile Signals

Simple (binary and/or directional) tactile signals are already commonplace in the form of mobile phone vibrotactile alerts for incoming calls; these are useful in many contexts where auditory signals are socially undesirable. Use of spatially distributed tactile signals has also been shown to speed up orientation of spatial attention, with a potential to aid in situational awareness [9], [90]. While signal complexity can be viewed as a continuum (defined either by information capacity in individual signals or by the number of uniquely recognizable signals achievable in a set), we are here defining simple signals as sitting at the far end of this continuum.

Value: The research to date suggests that simple signals are preferable to complex signals when 1) they are all that can be reliably detected, due to limitations of either hardware or context of use (e.g., when a cell phone is sitting in a pocket, details of the signal will be harder to make out), 2) only limited information need be conveyed, or 3) a strong, fast, and accurate user response is needed. By analogy, if visual attention is to be captured by a flashing light, response will be enhanced if that type of stimulus is only used for one event, rather than many different events indicated by variants in flash frequency or color, thus engaging a cognitive component in the response.

Choice of Hardware: For existing vibrotactile display hardware, there is a direct tradeoff between signal richness (potential complexity) and strength, particularly for power-starved mobile applications. For example, solenoid vibration is capable of much stronger stimuli, which can be noticed through clothing, as compared to more expressive configurations of piezo actuators; but it cannot create as many distinguishable signals, even when touched directly. Simple signals are also the more feasible option for less sensitive, nonglabrous skin delivery sites.

Abstract Communication and Information Display: Haptic Icons

The idea of using tactile signals to display abstractions has roots in communication aids for the blind, with the Optacon particularly notable [58]. A recent review of this application space can be found in [96], backed by reviews of relevant aspects of tactile psychophysics [35], [49], [77]. Abstract tactile information transmission has centered on haptic icons or their equivalent: brief informative haptic stimuli (usually vibratory) to which information has been attached.

Symbolic or Abstract: Haptic signals can be based on metaphorically derived symbols or more arbitrarily assigned associations. The likely pros and cons are fairly obvious. Symbolic notations intuitively seem easier to learn and remember, but there are obstacles to using this approach for large but usable sets of icons, particularly when the rendering palette is limited (imagine how well symbolic graphics would work using a few grayscale pixels to cover all possibilities). These challenges include independent control of signal salience and of perceptual spacing (some signals might feel very similar, others quite different, with no semantic pattern); and the fact that individuals are rarely consistent in their interpretations anyway—so one notation will not work for everyone. Both of these problems are handled relatively easily when the need for semiotic connection is dropped, e.g., using a process of *perceptual optimization* on a proposed signal set (see [61] and the discussion later in this article).

One approach to increasing the controllability of the semantic approach is to carefully ascertain a set of basic primitives with the goal of then using them across contexts in a variety of situations [97]. Another is for designers to create codes by drawing on an existing user knowledge base [13], [16]. Alternatively, we see that users are well able to create their own semantic mappings when given the means in both emotive [12], [18], [33] and informative [28] examples. In the last, we see what may be a cue for how to join the two approaches. A designer inflicted completely arbitrary links on his subjects, then discovered post hoc that most users created their own semantic mnemonics when learning the links, and typically found these personally derived interpretations just as logical (and learned them as well) as when they chose the stimulus-meaning associations themselves. That is: perhaps we can make anything behave as a semiotic link.

Learning Haptically Represented Abstractions: Regardless of the approach used to construct a stimulus-meaning link, in deploying the haptic channel for this kind of abstracted information transmission, we are asking individuals to use their touch sense in a manner they do not encounter in the natural world. Psychophysical evidence for tactile acuity with respect to this kind of information transmission is summarized in [63]. There is some neural evidence of brain plasticity for users asked to pick up this skill after early childhood [35], [43].

What learning techniques will best exploit this plasticity? Taking encouragement from human ability to learn Braille after childhood [39] and guidance from how it is taught, we note that a first step is generally to develop the learner's tactual acuity. Barraga and Errin describe a five-step process that moves from simple to complex, beginning with awareness and attention to tactile details, moving through recognition of structure and shape, part-to-whole relationships, then abstracted graphic representations and finally the learning of Braille symbols [5]. Immersion in rich and guided haptic experiences are the key in early stages [10], with Braille labeling introduced later [5].

Individual Differences: There appears to be significant individual variation in tactile acuity and ability to learn abstract associations, including both hyperacuity [21] and our own informal observations of a "haptically challenged" group among our typical experiment recruits. We do not yet know whether this range arises through basic perceptual function or learned cognition, and if the latter, what the indicators could be. Differences in how individuals organize their perceptual space have also been noted, with strong dimensions being held in common but different weaker dimensions employed differently [45]. Both types of difference (ability and organization) have implications on the widespread introduction of haptic information displays. An important area of future work is to better attribute the causes of both poor and exemplary haptic perceptual function, and to ascertain whether training and awareness can improve the former [66].

Identifying the Perceptual Dimensions of a Device Display Space: To create a set of learnable haptic icons, there are two linked challenges. One of these is creating learnable stimulus-meaning associations. Techniques for this are today largely ad hoc. The other is to ensure that the stimuli in the set are perceptually discernable, and furthermore to understand people's preference for organizing them, for later leverage in choosing appropriate patterns for association. For this, methods are more straightforward and there already exist the beginnings of a practical cataloging of the dimensionality and recognizable resolution available for various types of display hardware [13], [100], [103]. The current status on dimensionality that has been found for various types of stimuli and display hardware is summarized in [63].

Here, we will mention the one systematic tool of which we are aware, which uses Multidimensional Scaling (MDS) to "perceptually optimize" a group of stimuli. In a 20–60-min session (depending on the set size), a few users can provide enough dissimilarity data about a stimulus set to reliably create a map that reveals the dimensions along which the subjects perceive the stimuli relative to one another [61], [78], [100]. This map can be used to 1) guide iterative revision of the stimulus set until a renewed map indicates that the desired perceptual spacing (not too close or too different) has been achieved [16], [61]; and 2) choose a subset of stimuli for actual use in an application, again according to their desired perceptual organization and spacing. This method can be used both for independent creation of stimuli intended for arbitrary mapping to meanings, and for adjustment of a prototype set of representational icons whose meanings are chosen a priori [16].

Learning Stimulus-Meaning Associations: Glossing over the current sketchy state of affairs on creating learnable stimulus-meaning associations, the next step is for users to learn the associations. Because learning generally works best when information is absorbed from different sources (observed for tactile stimuli as well, e.g., [67]), a multisensory reinforcement learning process is probably advantageous even to learn a stimulus that might later be invoked purely through the haptic channel.

In efforts to date, users have already demonstrated a good ability to learn associations that are metaphorically matched by the designer [13], [16], [97], deliberately arbitrary [28], [29], or chosen by the user. In these instances, training took the form of repeated exposure/testing cycles of stimulus-meaning pairs until a given performance is demonstrated. We have also taken a further step of testing and continuing to optimize the icons under realistic environmental stress testing, adjusting the stimuli for relative distinctiveness and salience as needed. For example, in some circumstances, a controlled degradation in noticing performance is desired on response to workload, with some important icons still being noticed but less critical ones washing out when more urgent tasks are in play [17].

Expressive Control

"Expressive" refers to the quality or power of expressing an attitude, emotion, or other communicative information. Based on how we use touch in the real world, physicality seems a completely natural, indeed essential property for control tasks requiring emotiveness or precision, and in particular, both at once. We propose some heuristics and a brief summary of haptic potential in this realm.

Expressive Capacity

We use this term to broadly describe the richness of a communication channel for any purpose: its dimensionality, continuousness, the degree of control it affords the user, and the ease and naturalness with which desired acts can be completed [61]. This can refer both to tools that support artistic or interpersonal communication, i.e., emotional expression; and more prosaically, sheer information capacity. This can be specifically articulated as:

a) *Density*: number of bits of information that can be transmitted
b) *Controllability*: accuracy of conveyance (expression by sender, transmission, and interpretation by recipient)
c) *Directness*: direct versus encoded nature of the required actions (in analogy to direct-manipulation versus command-line interfaces)
d) *Responsiveness*: the immediate confirmatory and/or aesthetic feedback to the user
e) *Emotiveness*: the number, range, and subtlety of emotions that can be expressed.

By this measure, a computer keyboard is highly expressive on the first two counts but fails miserably in the third and fourth. The fifth is tricky, for the product of typing (the printed word) can be highly emotive in every way, both visually (ask a typesetter) and semantically. However, the *act* of typing is not particularly emotive. This raises the interesting question of whether an input device should be classified as expressive (based on its output) if using it doesn't *feel* expressive.

Role for Haptics

An ungrounded gestural interface works well for purely emotive control (low controllability). A keyboard is hard to beat when you wish to indirectly but exactly specify the greatest possible range of actions (high controllability). Physicality seems key when you need to do both at the same time. For example, in the highly studied topic of computer music controllers, many argue that the resistance and feedback of forces or vibrations are essential to controllability [20], [83], [104]. This is further linked to a consistency or closing of the control loop—a mechanical interaction between the subject and the sound source [37], [59]. However, computer-controlled grounded forces bring constraints, such as tethering and a loss of workspace, weight, motors and electrical power, a lack of generality in the control actions and handles that can be used, and a need for extremely tight synchronization between action and sound [7].

Some recent resources give guidance in how to accomplish this, from the standpoint of both the fundamental interactions themselves and their mechatronic implementation [11], [20], [62], [74]. Recent literature applying haptics to both music control and other expressive uses—ranging from the feel of a bristled paintbrush to gaming, control of under-actuated systems, and surgical simulation—is reviewed in [63]. A common feature is strong individuation of instrument to application, i.e., type of music to be created and the gestures employed. These are not general-purpose devices.

Haptic Effect

Affective design addresses the subjective emotional response to and relationship between users and interfaces. Although related, it is distinct from and more personal than expressive control: the latter is about achieving a desired result, although this does include the satisfaction and aesthetics of doing so. In the last decade, subjective response has been recognized as an important, if difficult-to-quantify aspect of everyday interfaces that impacts stress and usability [70]. It also forms the basis of a new, sophisticated type of interface based on *affective computing* [80], where the computer sensors and displays are used to determine and elicit particular emotional experiences from the user.

Haptic affective design has not received a lot of attention to date, despite recognition of the crucial role of touch in human communication and development [63]. Here, we mention two potential roles for effect in haptic design.

Design for Feel

Consider the direct affective response that feeling produces on the user: haptically speaking, what feels good, bad, or neutral? To what extent is this shaped by the task at hand? Is it consistent across people and does it impact performance? Preliminary efforts have explored mechanisms for measuring haptically induced effect (with a combination of biometric study and self-reports), is able to find some consistency in response, and suggests that haptic preference is not always linked to superior performance—i.e., sometimes people prefer controls that don't particularly aid in their task [93]. Eventually, this line of research will deliver heuristics that can guide interface aspects such as the choice of feel for a given control action. For now, the best practice is to routinely include subjective questionnaires in any performance-oriented user test during the design process, and consider this response in design iterations.

More broadly, we need clearer metrics to establish how important it is to get this right. The cost of negative affective response to an interface (whether the reaction is to ugly graphics, sound, or feel) is subtle and probably cumulative. One would expect the impact to be indirect but potentially far reaching, e.g., heightened tension and a lack of well being.

Emotional Communication

How can a haptic channel support human emotional communication? As noted in [92], current collaborative systems demand explicit communication—symbolic, focused, and overt, with an emphasis on transferring information in support of a goal. The overall situation hasn't changed much in the intervening decade, despite many experimental efforts aimed at understanding nonverbal human communication and attempting to support it remotely.

Mediated social touch is "the ability of one actor to touch another over a distance by means of tactile or kinesthetic feedback technology" (for an excellent review, see [42]). A number of examples using haptics have been explored, using a variety of direct force connections or tactile taps and with purposes ranging from emotional connectedness to therapy and ambient communication (summarized in [63]). They are provocative and insightful, but together demonstrate that we need a more systematic investigation of how, exactly, we communicate emotion through touch alone. Early evidence is that we can do so [4], [88] in at least simplified contexts. In another work, we are building a touch sensing-and-display platform to study this in a less-constrained environment [107].

Haptic Interaction Design Practices

There is a wealth of information on the best practices for UI design—textbooks [8], [24], [91], courses, conferences, and journals. There is also a growing literature on the principles for multimodal interface design, which is relevant here [75], [86]. Nevertheless, what is special about the process of designing haptics into interfaces? Or, even better, designing the interface itself around the idea of physical interaction?

Technocentric Versus User Centric Design

Because this article appears in a robotics magazine, it is a good guess that most readers have a technical background and are highly skilled at making machines do things. This can be a big problem when it comes to creating systems that work well for people, for a couple of reasons. The comments that follow are in no way limited to the design of haptic interfaces. But we are particularly vulnerable: haptic feedback started with robotics, and arises out of a culture of respect (reverence?) for complexity and automation. Although for nearly a decade now it's been possible to design haptic applications without building your own device, the ones you can buy mostly don't do quite what you need them to, and the technology is young and demanding enough that it still attracts practitioners of a tinkering mentality.

Have Need, Then Seek Technology: If you're looking for an application that will show off your device's special features, you 1) might have a fruitless search or 2) could find a good match, but then fail to do a good job of integrating it. Human problems are usually better solved by looking closely at the need, then surveying technologies to find the best match. This isn't much help if you're the engineer and have spent a lot of time building a cool gadget. It's important to watch and listen to people and notice where they struggle, and to hold an open mind. Perhaps your original solution isn't the right one, but the problem is real and understanding it will guide you to a different and better one.

Multidisciplinary Teamwork: Cultivate friends and associates who aren't engineers. By far the most productive design teams we've worked with, whether professionals or students, contain technologists, interaction designers, end users (including those with special needs or profession), and artists, freely and respectfully sharing ideas and possibilities. The most effective individual designers are empowered to observe, envision, and build—all within one brain and set of hands. So, leave your own comfort zone and learn to do what your partners are doing too.

Define Requirements in Solution-Independent Terms: When you do identify what seems to be a good problem-technology match, don't just jump in. This means studying the people you hope to help, and what they do without the proposed fixes. Talk to them, understanding that they won't always be able to articulate problems or envision hugely different solutions. Identify what's needed in *solution-independent terms*. Then, and only then, is it time to formulate specific designs with their enabling technology and begin to refine them.

Designing for an Unfamiliar Modality

Haptic design does differ in a significant way from visual and auditory design, in that most users will be initially unfamiliar with most possible uses for haptic technology. This is difficult enough when you're trying to simulate reality in some way, but becomes even harder when you create sensations or interactions that don't occur at all in the natural world.

Lost in Translation: It is difficult to predict how a programmed sensation will feel or whether an interaction will help until you build it and compare it against other possibilities. This is partly a matter of unmodeled device dynamics and partly of uncatalogued perceptual sensitivity. When will a sensation be masked or attenuated by another? Design iteration needs to include feedback from humans (perceptual questions) and sample end users (interaction questions).

Difficulty of Status Quo Comparisons: We often wish to know whether a haptic version or augmentation of a traditional visual interface helps people do something better, and seek a way to compare them. However, it can be difficult or pointless to create comparable versions. They are likely to be different in *many* ways, and so you must choose between a highly controlled comparison where one version is not optimally configured, or a poorly controlled comparison where it's hard to identify causal factors. We believe the most informative compromise is often to compare the best-of-breed versions and focus on collecting and analyzing rich observational data, in contrast to a hypothesis-testing approach, which emphasizes quantitative performance measures and statistical differences.

Evaluation in the Middle of the Learning Curve: The playing field isn't always level. For example, our subjects have been using vision for the kinds of tasks we test since early childhood, and they've been using the tactile version for perhaps a 3–30-min training period. It can be difficult to determine whether an innovation has intrinsic value, or extrapolate where it will go with experience. Longitudinal studies where subjects have more opportunity to become familiar with the use of haptics are expensive but clearly necessary.

Haptic Representations and Verbalizing Sensations: People aren't accustomed to processing haptic representations of abstractions, and they don't have a vocabulary to describe or help them remember detailed haptic distinctions the way they do for sounds and colors. As designers, we don't have a clear idea of the design dimensions. We've made a small start at correcting this [94].

Importance of Rapid Prototyping and Haptic Representations

Regardless of detail, a well-recognized principle of prototyping is to iterate at increasing levels of detail, whether creating a piece of software or a mechanical linkage. You don't start by building a refined, feature-complete instantiation of your vision, because it is likely to be wrong in many ways; and then you will have wasted a lot of effort. It is far more expensive to make changes late in the process when details become

Haptic interactions employ mechanical, programmed physical devices.

rigidified than in early, conceptual stages. For UIs, the truth of this maxim grows. Although there are some trustworthy heuristics, it is difficult to predict user response to any kind of novelty—modeling, simulation, or established rules are sparse and can be difficult to apply. For haptic UIs, the unfamiliarity and the combination of hardware and software design further amplifies this.

Minimalism: Prototyping UIs is an activity that lies somewhere between art, psychology, and science. Little can be described or left to the imagination, since users don't have a useful reference point. However, when prototypes are too high-fidelity early in the design cycle, they can appear to be finalized to a user; who will then be less likely to challenge or suggest modifications.

Modular Prototyping: The primary objective of a prototype is to get your design question answered with as little effort as possible, starting with big ones and proceeding to more detailed ones. It can then be discarded when you move on. If you have an engineering feasibility question, then implement exactly the degree of functionality needed to test that. If you need to figure out if a physical configuration is going to work for a user, then a nonactuated mockup might allow you to get this feedback from a user for a lot less work than a functional model. If you need to test look-and-feel or aesthetics, a conceptual or even a graphical rendering could be sufficient.

Later in the process, it makes sense to prototype multiple aspects together. It's more expensive but by now major directions are confirmed, and risk is gradually being reduced. You'll continue to make new discoveries as more of the system comes online, and you are able to observe real users interacting with increasingly realistic and functional mockups. This modularity is illustrated in the first case study presented later.

Brainstorming and Multiple Approaches: Pursuing a single path to a design goal is unlikely to give the best result. Brainstorming (the wild, absurd kind) helps to generate creative, far-flung approaches which when recombined, toned down, and refined can open up new directions. Whenever possible, advance two or three different paths that are as different as possible. In the end, you'll likely combine elements of different approaches, and you'll have more understanding of the design landscape.

Tools: The principal danger of tools is their introduction of an insidious obstacle to innovation in alternate directions. Having a choice of tools and being aware of their constraints is helpful.

Triangulation in Prototype Creation and Evaluation: Each prototype is built to be evaluated in some way, whether mechanically or in terms of comprehensibility or aesthetics. Any kind of evaluation is flawed, in part because you're only prototyping and observing part of the whole experience. *Triangulation* refers to coming at each evaluative point from multiple directions, using techniques whose strengths and weaknesses complement one another. For example, performance-based and observational evaluations provide different views. For more on user evaluation, see an introductory HCI textbook (e.g., [8], [24], [91]).

Prototyping Things That Can't Be Built: As for any novel technology, to advance we often need to prototype the future. Today's hardware limits us, but if we can show real value for a technology we can't yet build, this can inspire development effort in that direction. For example, our group has put tactile displays into hand-held devices that cannot yet be built with sufficient compactness and power efficiency to actually be untethered. However, we won't know if it's worth finding a way to make this technical advance or be ready for it when it comes, if we haven't by then found a way to use it effectively.

Some Ideas for Getting Started

You have your real human problem, a technology that seems like it should help, and you're prepared to prototype. How do you start?

Each design problem is unique, and we're not at the point of recipes. Nevertheless, we can suggest some ways to get going, which may even end up as useful design approaches.

Use of Metaphor

When an information or control task has roots in predigital interactions, exploiting these roots by building metaphorical interactions around them can aid control and make it comprehensible. An example of this is introduced in the first case study, which describes a mediating virtual physical metaphor for interacting with media. The haptic representation is not of the media itself but of a virtual tool which has similarities to one that users might have once used in the real world [62], [89].

Navigating Modes

Haptic feedback is often proposed as a solution for *modal interfaces* in which the interface can be in different states, and a command thus means different things depending on the state. Problems arise with modal interfaces when the current state is not evident, or when it's hard to move between them. A haptic display (say, a knob with an embedded liquid crystal display) has possibilities here because unlike a physical knob, it can be reprogrammed appropriately for the current mode, just like the graphical display. However, when the graphical display goes away—or the user can't look at it for a while—then the haptic display must be able to transparently indicate the mode. The state of the art in our current hardware is point-based interaction for force feedback. This means that usually you have to explore an environment serially to deduce the state. This is undesirable, and you might inadvertently alter the state in the process. How can we get around this?

One approach is to redefine the interaction in a manner that either gets rid of modes altogether or allows the user's active, deliberate motions to alter or navigate through them in an intuitive way. At the same time, the interface can supply ongoing physical feedback about the state, without requiring continual system interrogation. Physical metaphor is a good way to enter into this idea, because it is how real hand-held tools work: e.g., you might shift the position of a tool in your hand or switch tools entirely (deliberate physical act) and then continue to receive feedback through the shape of the tool in your hand and the sensations transmitted during its use (think about how different writing and cutting implements feel, in terms of shape, heft, and transmitted forces and vibrations). It is hard to change the shape of a handle, but you might be able to change its virtual weight or center of inertia and certainly the vibrations.

Modal Continuums: Discrete and Continuous Control

We think of interface modes as being discrete states, but sometimes this is an artificial construct, and, in fact, the desired control shifts along a continuum. Again using the digital media example, observe how when traversing a media stream we move between discrete and continuous forms of the material, its content, and aggregations. Video is a succession of frames—discrete when played slowly but merged into a fluid when sped up. Spinning the virtual video reel of the first case study allows one to move seamlessly between these states: the ticking of individual frames speeds into a texture while the frame rate fuses visually. A collection of voice mail messages, music tracks, or cable TV channels are discrete objects; when played, individual items are continuous streams. If the set is represented in the right way, you can skim over the discrete items themselves like a texture, feeling for the variation that indicates the item property you're looking for.

Design Case Studies

We conclude with a pair of case studies that illustrate ways in which haptic feedback can be explicitly designed for an application context, chosen to span a broad space of application areas and a variety of principled design mechanisms. For authenticity and detail as well as brevity and focus, they are chosen from the authors' own experience.

Force Feedback Knob: Continuous and Discrete Hand-Held Media Control

Along with digitization of once-tangible tasks and microcomputers everywhere, comes the frequent necessity of managing information or controlling systems through very simple input devices. When hapticized, this generally comes down to knobs and sliders.

In this first example, we relate key points of a design sequence that relied on metaphor to create generalizable but experience-grounded interactions for a hand-held media controller [64], [65], [89], beginning with some relevant principles and observations. The first stages of this project were performed at Interval Research Corp. (Palo Alto, CA) during 1998–1999 by a design team lead by the first author. Later stages were conducted as student projects at University of British Columbia, Canada. This case also illustrates the modular prototyping principle described in the previous section. Starting from the ideas of metaphor-based design and discrete/continuous media modes, we set out to build a hand-held home media controller that would leverage the utility of modal interaction for different media in a consistent way, while making the state transparently clear.

Inspiring Metaphor: We tried a lot of metaphors! And ended up using several. One which felt good and aided navigation was a virtual "clutch" through which the user interacted with a heavy "reel" of film that runs on the computer screen as the reel spins (Figure 1). The inspiration for the bit of applied tangibility used here came from discussions with videographers who missed some aspects of traditional mechanisms for handling celluloid film. It allows a far more fluid handling of the information than cursor clicks or stop/start buttons.

Technical Grounding: A technical path for this was suggested by *tangible interfaces*, where tagged arbitrary objects (e.g., using radio frequency or RFID) can be used to issue commands to a computer [101]. Observing that tagged objects are well suited for issuing digital commands but not for exerting continuous control, we combined the two through the principle of *tagged handles* [64], [65].

Prototype-Driven Design Steps

Figure 2 illustrates several successive prototypes in an iterative conceptual and engineering evolution. In this process, exploration of the prototypes themselves drove further designs, and there was an emphasis on lightweight prototyping where possible. These began with an engineering exercise, shared informally with users, to see whether the combination of discrete (tagged handle) and continuous (force feedback knob) would be compelling [see the prototype in Figure 2(a)]. Each of the handles contained a unique RFID tag, which when installed on the force feedback knob caused the system to browse (and give appropriate force feedback for) a particular kind of content or functionality—e.g., a particular music track, selection of radio versus recorded content, or volume versus navigational control. No attempt was made at usability—for example, the handles did not suggest their function.

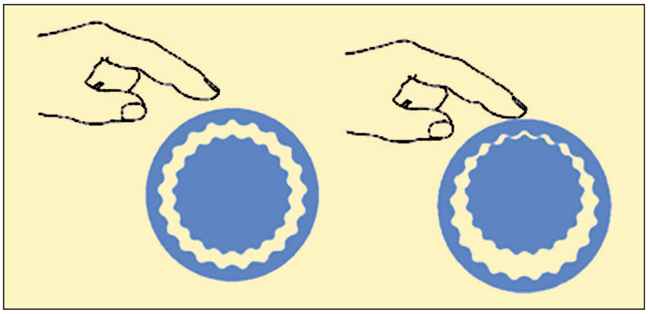

Figure 1. *Virtual clutch metaphor for the force feedback media controller. The knob is equipped with a crude pressure sensor. When the user presses down on the actual knob (which is associated with the outer wheel in this figure), the heavy inner wheel (virtual) is engaged and can be spun up. When the actual knob is released, the inner wheel continues with its imparted inertia. The video displayed on the screen is linked to the rotational speed of the virtual inner wheel. The bumps displayed here correspond to frames and are haptically rendered as small detents that fuse into a texture as the speed increases.*

The prototype in Figure 2(b) is an example of the many ideas explored mostly at a conceptual level in Figure 3. It is a nonfunctioning prototype showing one way that discrete handles (inspired by a charm bracelet) could informatively indicate their function and solve the practical problem of getting lost—the handle is selected from a wheel instead of being picked up and attached. Sadly, these protruding little handles would take a finger off when it rotated under active control, and several more nonfunctional prototypes (not shown) led to the next step.

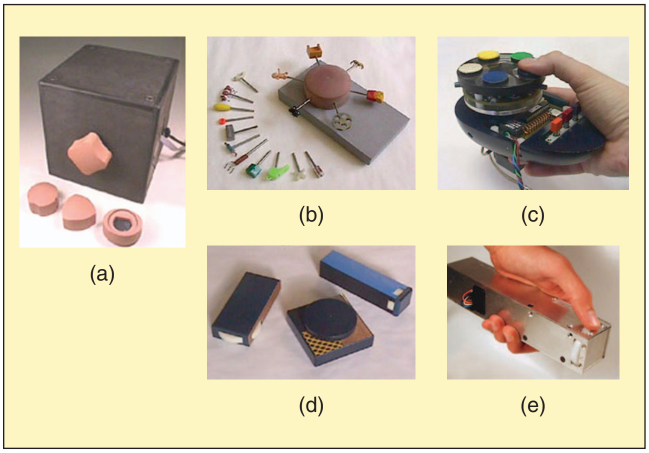

Figure 2. *Representative haptic media controller design iterations: (a) The initial tagged handles engineering concept prototype. (b) A representative conceptual prototype. (c) A later technical prototype (oversized). (d) A set of nonfunctional concept prototypes that addresses the problems of (c). (e) Another engineering mockup.*

The prototype in Figure 2(c) is a fully functional implementation of a safer variant of the same idea—handles are replaced by texturally marked buttons on a rotating wheel mounted on a hand-held base. In using this mockup, we discovered a problem of disorientation. When the face rotated, the buttons moved, and they were hard to find again; spatial constancy turns out to be critical. The next refinement [Figure 2(d)] inverted this idea. A four-sided object with texturally marked sides and an active thumbwheel knows which face is active by measuring where the thumbwheel was pressed from and changing the function and feedback of the continuous interaction accordingly—e.g., turning from one face could change volume, and turning from another could select channel. Finally, Figure 2(e) is another engineering prototype of this final idea [64].

In Summary

This case study showed a prototype-dominated process, where user feedback was obtained informally at each stage. The use of varied, focused, and stage-appropriate prototypes allowed us to identify key strengths and weaknesses with minimal effort. This example did not make use of extensive, formally controlled user studies for feedback on the prototypes because the concept clearly had many bugs to be worked out before we even reached that stage. However, it was inspired and informed by parallel efforts at the host company, consisting of extensive ethnographic studies of target user groups in their uses of home media, and interviews focusing on their difficulties with currently available models. That is, the user-centered component was upfront observation, and the next step would have been a usability study . . . if the host company hadn't vaporized in the 2000 tech bust.

Figure 3. *Early prototype for the hand-held media controller project. Found objects and state-of-art examples, Lego + rubber band transmissions, whimsical and serious nonfunctioning concepts, and narrowly targeted functional engineered prototypes.*

Vibrotactile Background Signals

Our second example, in contrast, is heavy on the user studies. Its goal was a first deployment of a set of haptic icons in an application concept. It began with devising an initial set of icons using a symbolic approach based on metaphors thought to intuitively represent the concepts being represented. The icon set was then systematically refined in an iterative, user-centered process mentioned previously (see the "Tactile Signaling in Multitasking Environments" section) and culminating in an *observational user study.* These steps are more fully described in [16] and [17], and we summarize some key points here.

Application: When noncolocated and collaborating users wish to jointly modify a shared object displayed on their local screens—whether a text document, a computer-aided design drawing or a Photoshop file—current technology (e.g., virtual network protocol or VNC) allows only one of them to control the cursor at a time. Somehow, they need to negotiate turn-taking, but in the absence of the nonverbal cues that are so important in colocated situations. (Our own guess is that even the usual nonverbal cues available in colocated meetings could use help too. Could tactile cues discretely remind someone who's impervious to coughs, raised hands, and squirming, that it's really time to stop monopolizing the floor?)

We began with the proposition that tactile feedback could provide a background awareness of others' wish to participate. It could indicate both turn-request queue and urgency of items in the queue, in a less intrusive manner than visual or auditory methods could support—because the latter were also being used in the collaborative task. We further wondered if the ability to make a request gently or urgently would support more equitable control sharing. A quiet or shy team member might be more comfortable asking for control "whenever you're ready," as opposed to "right now!." It was problematic for visual or auditory protocols to support this. Requests not dealt with right away couldn't easily persist, because they'd either be in the way or forgotten.

The only way we could test this idea (which we hoped was representative of a whole class of applications) was to build up a set of icons and try it out on users in a realistic situation.

Experiment Paradigm and Display Hardware: The climactic observational study involved groups of four friends who were placed out of direct eye- and earshot of one another (Figure 4) and given voice links and a shared screen view of a common application (a furniture-layout task using Visio). They received tactile feedback through modified tactile mice (Logitech IFeel; Figure 5). Although more expressive displays were available, we wanted to see how far we could get with commodity hardware. Groups performed the room layout task three times: with only tactile mediation, with only visual mediation (following state-of-art visual protocols), and using both modalities. Each member was given responsibility for a subset of the criteria that had to be followed in the solution, and the group collectively got a bonus if they did particularly well. Their interactions were closely monitored.

Haptic feedback has been shown to have value in the training of sensorimotor tasks.

Protocol and Initial Icon Creation: With this scenario in mind, we designed the turn-taking protocol and the initial set of haptic stimuli that would support it, as well as the analogous visual signals. In essence, the protocol recognized three classes of users—those in control, those waiting for control, and those just observing; two types of requests—urgent and gentle; and two types of events—an urgent or a gentle request and a self-removal from the queue. Seven icons were needed to display the current context as relevant to a given user. For example, the user who was in control would experience a different signal than one who was in the queue. The haptic stimuli which were eventually used, are shown in Figure 6. We

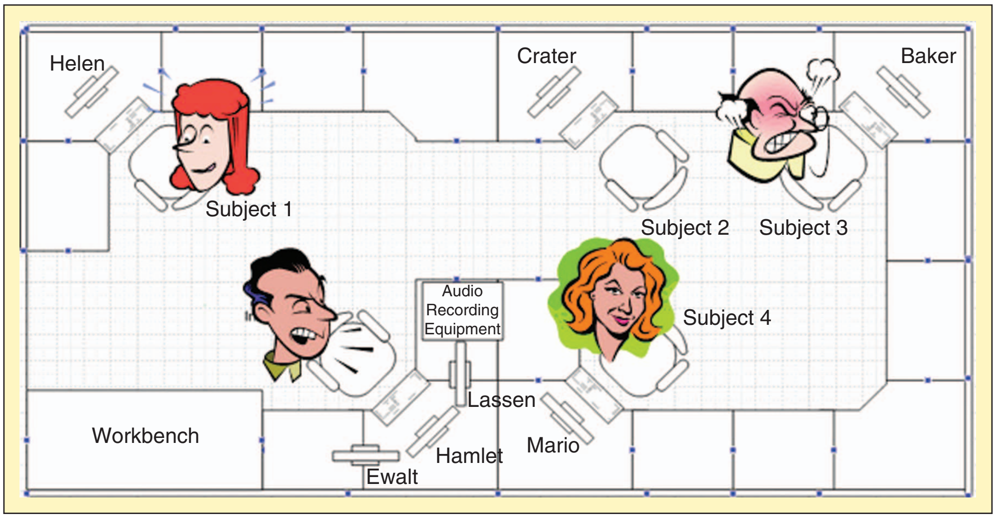

Figure 4. *Experimental setup for the observational study of the turn-taking protocol. The four group members were placed out of direct eyeshot and wore noise-canceling headphones; all vocal communications occurred through a sound system.*

Figure 5. *Vibrotactile mouse used to display the haptic icons used in the turn-taking protocol. Two buttons were added to the side to enable special protocol features; such buttons were available in other mouses at the time but not the vibrotactile one.*

The presence of the visual and auditory stimuli can significantly modify the user's interpretation of what they feel.

used a metaphor-based design on the assumption that it would make this small set easier to learn. For example, the change of control states were suggestive of the be-BEEP, BE-beep of the common auditory cue indicating the insertion or removal of a hardware device from your computer.

Process: User-Focused Icon Set Refinement and In Situ Observation

We were too experienced with haptic icons to think we were ready for prime time, though. Would users actually be able to learn them? Would they be confused with one another? Was their salience correctly adjusted? We thus commenced on a multistep refinement process. The initial icon set design described above was Step I (we're currently working on alternatives to its fairly ad hoc nature).

In Step II, we perceptually adjusted the icon set using the MDS technique mentioned previously, testing the most likely candidates along with a lot of others. A few iterations of this served to ensure that all the icons in the set were well distributed within the engineering design space.

In Step III, we "stress-tested" the icon set in realistic conditions, by requiring subjects to learn associations, then abstractly simulating various aspects of the anticipated workload (with appropriate visual and auditory load), and examining how icon detection and identification degraded. We wanted some icons to be less detectable under workload, while others should always get through. For example, an *in control* user should always perceive and recognize an urgent request, but while concentrating hard, he shouldn't be bothered with a gentle request—that was the whole point of the urgency-based protocol. Following this test, we adjusted some of the signals even more to get the desired salience patterns. Subjects learned the seven mediating icons easily in three minutes and maintained 97% accuracy of identification under substantial multimodal workload.

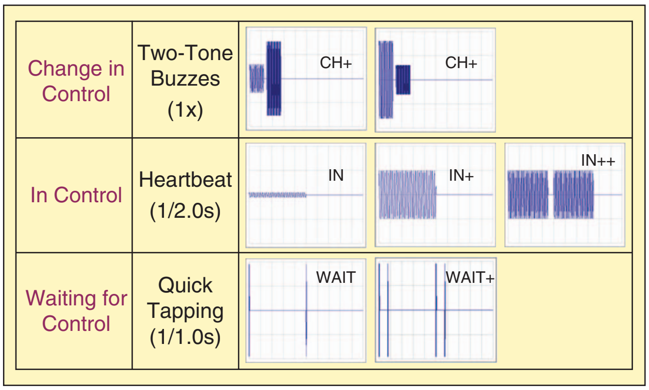

Figure 6. *Final set of haptic icons used in the turn-taking protocol.*

Unfortunately, we did not then return to Step II to readjust their perceptual spacing; next time we will! The salience adjustment did, we later learned, make some pairs harder to distinguish.

Finally, in Step IV, we mounted the group observational study, and learned quite a lot (read the article). Through a combination of performance and subjective measures we did confirm that the haptic signals were utilized in a graded (i.e., appropriate) way, and collaboration dynamics seemed to be positively affected in comparison to the visual cue case. Users, however, preferred having *both* visual and haptic cues available to them.

In Summary

This case exemplifies a quite user-intensive design process. The hardware itself was simple, but what we did with it would fail or succeed based on subtle details, and this could only be determined by trying it out while watching closely. The final endeavor was an observational rather than tightly controlled, performance-oriented study, out of a combination of necessity and design. Because each session was a lot of work, we could only run four groups of varied background, and thus there wasn't enough data to give statistical results. However, by observing and logging everything and following up with detailed interviews (and a second set of interviews a month later after looking over the data) we obtained a great deal of complex and nuanced feedback on the strengths and weaknesses of the approach. Given that there are many ways to implement this general concept, observational data were more valuable at this stage than hard performance data.

Summary

In this second part of our series, we have introduced the concept of and argued the need for explicit, user-centered interaction design for applications using haptic interfaces. We elaborated on a number of potential interface roles where haptic feedback is well suited to provide value, on the basis of the technology's alignment with human capabilities and modern needs, and we suggested some high-level principles to be followed and the pitfalls to be avoided during the application design process. Finally, we illustrated these with two case studies, chosen for their different approaches to the interaction design process.

Readers who are interested in learning more should start by learning about HCI practices in general, through textbooks and courses. Many aspects of user-centered design practices apply here but are unfamiliar to the engineering world. A working knowledge of haptic perception is essential as well. Because this frontier is advancing so rapidly, simply following these articles in haptics conferences will get you far, as well as the survey material mentioned earlier.

In Part I, we introduced the haptic devices themselves, their construction, and operating principle and placed special emphasis on some simple display variants that can be constructed and employed with little special expertise. We hope that our comments in Part II, in tandem with the electromechanical design principles in Part I, will lower the barrier to entry for this exciting young field, and foment many new ideas—usable ones!

Haptic design is nearly always a multimodal design.

Keywords

Haptic interfaces, interaction design, ubiquitous computing, force feedback, tactile feedback, human computer interaction.

References

[1] R. J. Adams, D. Klowden, and B. Hannaford, "Virtual training for a manual assembly task," *Haptics-e*, vol. 2, no. 2, 2001.

[2] C. A. Avizzano, J. Solis, and M. Bergamasco, "Teaching to write Japanese characters using a haptic interface," presented at the *10th Symp. Haptic Interfaces for Virtual Environment and Teleoperator Systems*, Washington, DC, 2002.

[3] P. Bach-y-Rita, K. A. Kaczmarek, M. E. Tyler, and M. Garcia-Lara, "Form perception with a 49-point electrotactile stimulus array on the tongue: A technical note," *J. Rehab. Res. Dev.*, vol. 35, no. 4, pp. 427–430, 1998.

[4] J. N. Bailenson, N. Yee, S. Brave, D. Merget, and D. Koslow, "Virtual interpersonal touch: Expressing and recognizing emotions through haptic devices," *Hum. Comput. Interact.*, to be published.

[5] N. Barraga and J. Erin, *Visual Handicaps and Learning*. Austin, TX: ProEd, 1992.

[6] C. Basdogan, C. H. Ho, M. Srinivasan, and M. Slater, "An experimental study on the role of touch in shared virtual environments," *ACM Trans. Comput. Hum. Interact. (TOCHI: Special Issue on Human Computer Interaction and Collaborative Virtual Environments)*, vol. 7, no. 4, pp. 443–460, 2000.

[7] T. Beamish, K. MacLean, and S. Fels, "Manipulating music: Multimodal interaction for djs," in *Proc. ACM Conf. Human Factors in Computing Systems (CHI'04), CHI Lett.*, Vienna, Austria, 2004, vol. 6, no. 1, pp. 327–334.

[8] D. Benyon, P. Turner, and S. Benyon, *Designing Interactive Systems: People, Activities, Contexts, Technologies*. Reading, MA: Addison Wesley, 2005.

[9] P. Bertelson and B. de Gelder, "The psychology of crossmodal attention," *Crossmodal Space and Crossmodal Attention*, C. Spence and J. Driver, Eds. Oxford: Oxford University Press, 2004, pp. 141–177.

[10] S. J. Blake. (2007). *Teaching Reading in Braille*. [Online]. Available: http://blindness.growingstrong.org/ed/aa010401a.htm

[11] A. J. Bongers, "Interactivation: Towards an e-cology of people, our technological environment, and the arts," Proefschrift, Vrije Universiteit, Amsterdam, The Netherlands, 2006.

[12] S. Brave and A. Dahley, "Intouch: A medium for haptic interpersonal communication," in *Proc. Conf. Human Factors in Computing Systems (CHI'97)*, pp. 363–364.

[13] L. M. Brown, S. A. Brewster, and H. C. Purchase, "Multidimensional tactons for non-visual information presentation in mobile devices," in *Proc. 8th Conf. Human-Computer Interaction with Mobile Devices and Services*, Helsinki, Finland, 2006, pp. 231–238.

[14] D. B. Camarillo, T. M. Krummel, and J. K. Salisbury, "Robotic technology in surgery: Past, present, and future," *Am. J. Surg.*, vol. 188, no. 4A, pp. 2S–15S, 2004.

[15] C. S. Campbell, S. Zhai, K. W. May, and P. Maglo, "What you feel must be what you see: Adding tactile feedback to the trackpoint," in *Proc. Interact'99*, Edinburgh, UK, pp. 383–390.

[16] A. Chan, K. E. MacLean, and J. McGrenere, "Learning and identifying haptic icons under workload," in *Proc. 1st Joint Eurohaptics Conf. Symp. Haptic Interfaces for Virtual Environment and Teleoperator Systems (WHC 2005, IEEE-VR2005)*, Pisa, Italy, pp. 432–439.

[17] A. Chan, K. E. MacLean, and J. McGrenere, "Designing haptic icons to support collaborative turn-taking," *Int. J. Hum. Comput. Stud.*, to be published.

[18] A. Chang, S. O'Modhrain, R. Jacob, E. Gunther, and H. Ishii, "Comtouch: Design of a vibrotactile communication device," in *Proc. Conf. Designing Interactive Systems (DIS'02)*, London, pp. 312–320.

[19] S. Choi and H. Z. Tan, "An analysis of perceptual instability during haptic texture rendering," in *Proc. 10th Symp. Haptic Interfaces for Virtual Environment and Teleoperator Systems*, Los Alamitos, CA, 2002, pp. 129–136.

[20] P. Cook, "Principles for designing computer music controllers," in *Proc. Conf. New Interfaces for Musical Expression*, Seattle, Washington, 2001, pp. 1–4.

[21] J. C. Craig, "Vibrotactile pattern perception: Extraordinary observers," *Science*, vol. 196, no. 4288, pp. 450–452, 1977.

[22] S. Delp, P. Loan, C. Basdogan, and J. M. Rosen, "Surgical simulation: An emerging technology for emergency medical training," *Presence Teleoper. Virtual Environ.*, vol. 6, no. 2, pp. 147–159, 1997.

[23] D. DiFilippo and D. K. Pai, "Contact interaction with integrated audio and haptics," presented at *Int. Conf. Auditory Display (ICAD'00)*.

[24] A. Dix, J. E. Finlay, G. D. Abowd, and R. Beale, *Human-Computer Interaction*, 3rd ed. Englewood Cliffs, NJ: Prentice Hall, 2003.

[25] S. R. Dixon, C. D. Wickens, and J. S. McCarley, "On the independence of compliance and reliance: Are automation false alarms worse than misses?" *Hum. Factors*, vol. 49, no. 4, pp. 564–572, 2007.

[26] M. Enriquez, O. Afonin, B. Yager, and K. MacLean, "A pneumatic tactile notification system for the driving environment," in *Proc. Workshop on Perceptive User Interfaces (PUI'01)*, Orlando, FL, pp. 1–7.

[27] M. Enriquez and K. MacLean, "Impact of haptic warning signal reliability in a time-and-safety-critical task," in *Proc. 12th Annu. Symp. Haptic Interfaces for Virtual Environments and Teleoperator Systems (IEEE-VR2004)*, Chicago, IL, pp. 407–415.

[28] M. Enriquez and K. E. MacLean, "The role of choice in longitudinal recall of meaningful tactile signals," presented at the *IEEE Symp. Haptic Interfaces (HAPTICS 2008)*, Reno, NV.

[29] M. Enriquez, K. E. MacLean, and C. Chita, "Haptic phonemes: Basic building blocks of haptic communication," in *Proc. 8th Int. Conf. Multimodal Interfaces (ICMI 2006)*, Banff, Canada, pp. 302–309.

[30] M. O. Ernst and M. S. Banks, "Humans integrate visual and haptic information in a statistically optimal fashion," *Nature*, vol. 415, pp. 429–433, 2002.

[31] D. Feygin, M. Keehner, and F. Tendick, "Haptic guidance: Experimental evaluation of a haptic training method for a perceptual motor skill," in *Proc. 10th Symp. Haptic Interfaces for Virtual Environment and Teleoperator Systems*, Los Alamitos, CA, 2002, pp. 40–47.

[32] J. Fogarty, S. E. Hudson, C. G. Atkeson, D. Avrahami, J. Forlizzi, S. Kiesler, J. C. Lee, and J. Yang, "Predicting human interruptibility with sensors," *ACM Trans. Comput. Hum. Interact.*, vol. 12, no. 1, pp. 119–146, 2005.

[33] B. Fogg, L. D. Cutler, P. Arnold, and C. Eisbach, "Handjive: A device for interpersonal haptic entertainment," in *Proc. Conf. Human Factors in Computing Systems*, 1998, pp. 57–64.

[34] B. Forsyth and K. E. MacLean, "Predictive haptic guidance: Intelligent user assistance for the control of dynamic tasks," *IEEE Trans. Visualizat. Comput. Graph.*, vol. 12, no. 1, pp. 103–113, 2006.

[35] A. Gallace, H. Z. Tan, and C. Spence, "The body surface as a communication system: The state of the art after 50 years," *Presence Teleoper. Virtual Environ.*, vol. 16, no. 6, pp. 655–676, Dec. 2007.

[36] J. J. Gibson, "Observations on active touch," *Psychol. Rev.*, vol. 69, no. 6, pp. 477–490, 1962.

[37] B. Gillespie, "The touchback keyboard," in *Proc. Int. Computer Music Conf.*, San Jose, CA, 1992, pp. 447–448.
[38] B. Gillespie, S. O'Modhrain, C. P. Tang, and D. Zaretsky, "The virtual teacher," in *Proc. ASME Dynamic Systems and Control Division*, 1998, vol. 64, pp. 171–178.
[39] A. Grant, M. Thiagarajah, and K. Sathian, "Tactile perception in blind Braille readers: A psychophysical study of acuity and hyperacuity using gratings and dot patterns," *Percept. Psychophys*, vol. 62, no. 2, pp. 301–12, 2000.
[40] W. B. Griffin, W. R. Provancher, and M. R. Cutkosky, "Feedback strategies for shared control in dexterous telemanipulation," in *Proc. IEEE/RSJ Int. Conf. Intelligent Robots and Systems (IROS 2003)*, vol. 3, pp. 2791–2796.
[41] P. Griffiths and R. B. Gillespie, "Shared control between human and machine: Haptic display of automation during manual control of vehicle heading," in *Proc. 12th Int. Symp. Haptic Interfaces for Virtual Environment and Teleoperator Systems (HAPTICS'04)*, pp. 358–366.
[42] A. Haans and W. Ijsselsteijn, "Mediated social touch: A review of current research and future directions," *Virtual Reality*, vol. 9, no. 2, pp. 149–159, 2006.
[43] R. H. Hamilton and A. Pascual-Leone, "Cortical plasticity associated with Braille learning," *Trends Cogn. Sci.*, vol. 2, no. 5, pp. 168–174, 1998.
[44] V. Hayward, "A brief taxonomy of tactile illusions and demonstrations that can be done in a hardware store," *Brain Res. Bull. (Special Issue on Robotics and Neuroscience)*, to be published.
[45] M. Hollins, S. Bensmaïa, K. Karlof, and F. Young, "Individual differences in perceptual space for tactile textures: Evidence from multidimensional scaling," *Percept. Psychophys.*, vol. 62, no. 8, pp. 1534–1544, 2000.
[46] E. Horvitz, C. Kadie, T. Paek, and D. Hovel, "Models of attention in computing and communication: From principles to applications," *Commun. ACM*, vol. 46, no. 3, pp. 52–59, 2003.
[47] R. D. Howe and Y. Matsuoka, "Robotics for surgery," *Annu. Rev. Biomed. Eng.*, vol. 1, pp. 211–240, 1999.
[48] S. C. Jacobsen, F. M. Smith, D. K. Backman, and E. K. Iversen, "High performance, high dexterity, force reflective teleoperator II," in *Proc. ANS Topical Meeting on Robotics and Remote Systems*, Albuquerque, NM, 1991, pp. 1–10.
[49] L. Jones and N. Sarter, "Tactile displays: Guidance for their design and application," *Hum. Factors: J. Hum. Factors Ergon. Soc.*, to be published.
[50] L. A. Jones, M. Nakamura, and B. Lockyer, "Development of a tactile vest," in *Proc. Haptic Interfaces for Virtual Environment and Teleoperator Systems (HAPTICS'04)*, Chicago, IL, pp. 82–89.
[51] R. Kikuuwe and T. Yoshikawa, "Haptic display device with fingertip presser for motion/force teaching to human," in *Proc. IEEE Int. Conf. Robotics and Automation*, 2001, pp. 868–873.
[52] D. Kragic, P. Marayong, M. Li, A. M. Okamura, and G. D. Hager, "Human machine collaborative systems for microsurgical applications," *Int. J. Robot. Res.*, vol. 24, no. 9, pp. 731–742, 2005.
[53] H. I. Krebs, N. Hogan, M. L. Aisen, and B. T. Volpe, "Robot-aided neuro-rehabilitation," *IEEE Trans. Rehabil. Eng.*, vol. 6, no. 1, pp. 75–87, 1998.
[54] R. Leung, K. E. MacLean, M. B. Bertelsen, and M. Saubhasik, "Evaluation of haptically augmented touch screen GUI elements under cognitive load," presented at the *9th Int. Conf. Multimodal Interfaces (ICMI'07)*, pp. 374–381.
[55] D. J. Levitin, K. E. MacLean, M. V. Mathews, L. Chu, and E. R. Jensen, "The perception of cross-modal simultaneity," *Comput. Anticipat. Syst.*, vol. 517, pp. 323–329, 2000.
[56] M. Lin and M. Otaduy, *Haptic Rendering: Foundations, Algorithms and Applications*. Wellesley, MA: A.K. Peters, to be published.
[57] M. Lin and R. H. Taylor, "Spatial motion constraints in medical robot using virtual fixtures generated by anatomy," in *Proc. IEEE Int. Conf. Robotics Automation (ICRA'04)*, vol. 2, pp. 1270–1275.
[58] J. G. Linvill and J. C. Bliss, "A direct translation reading aid for the blind," in *Proc. IEEE*, 1966, vol. 54, pp. 40–51.
[59] A. Luciani, J. Florens, and N. Castagne, "From action to sound: A challenging perspective for haptics," in *Proc. 1st Joint Eurohaptics Conf. and Symp. Haptic Interfaces for Virtual Environment and Teleoperator Systems (WHC'05)*, pp. 592–595.
[60] J. Luk, J. Pasquero, S. Little, K. MacLean, V. Levesque, and V. Hayward, "A role for haptics in mobile interaction: Initial design using a handheld tactile display prototype," in *Proc. ACM Conf. Human Factors in Computing Systems (CHI'06)*, Montreal, Canada, vol. 8, no. 1, pp. 171–180.
[61] K. MacLean and M. Enriquez, "Perceptual design of haptic icons," presented at *Eurohaptics*, Dublin, Ireland, 2003.
[62] K. E. MacLean, "Designing with haptic feedback," in *Proc. IEEE Robotics and Automation (ICRA'00)*, San Francisco, CA, vol. 1, pp. 783–788.
[63] K. E. MacLean, "Haptics in the wild: Interaction design for everyday interfaces," *Reviews of Human Factors and Ergonomics*, M. Carswell, Ed. Santa Monica, CA: Human Factors and Ergonomics Society, to be published.
[64] K. E. MacLean, M. J. Shaver, and D. K. Pai, "Handheld haptics: A USB media controller with force sensing," in *Proc. IEEE VR2002 10th Symp. Haptic Interfaces for Virtual Environment and Teleoperator Systems*, Orlando, FL pp. 311–318.
[65] K. E. MacLean, S. S. Snibbe, and G. Levin, "Tagged handles: Merging discrete and continuous control," in *Proc. ACM Conf. Human Factors in Computing Systems (CHI'00), CHI Lett.*, The Hague, Netherlands, vol. 2, no. 1, pp. 225–232.
[66] S. Millar, *Reading by Touch*. London: Routledge, 1997.
[67] S. Millar, "Memory in touch," *Psicothema*, vol. 11, no. 4, pp. 747–767, 1999.
[68] T. Moran and P. Dourish, "Introduction to this special issue on context-aware computing," *Hum. Comput. Interact.*, vol. 16, no. 2, pp. 87–95, 2001.
[69] D. Morris, H. Tan, F. Barbagli, T. Chang, and K. Salisbury, "Haptic feedback enhances force skill learning," in *Proc. IEEE World Haptics*, 2007, pp. 21–26.
[70] D. A. Norman, *Emotional Design: Why We Love (or Hate) Everyday Things*. New York: Basic Books, 2003.
[71] D. A. Norman, "The next UI breakthrough, Part 2: Physicality," *ACM Interact.*, vol. 14, no. 3, pp. 44–45, 2007.
[72] S. S. Nudehi, R. Mukherjee, and M. Ghodoussi, "A shared-control approach to haptic interface design for minimally invasive telesurgical training," *IEEE Trans. Control Syst. Technol.*, vol. 13, no. 4, pp. 588–592, 2005.
[73] M. K. O'Malley, A. Gupta, M. Gen, and Y. Li, "Shared control in haptic systems for performance enhancement and training," *J. Dyn. Syst. Meas. Control*, vol. 128, no. 1, pp. 75–85, 2006.
[74] D. O'Sullivan and T. Igoe, *Physical Computing*. Boston, MA: Thomson Course Technology, 2004.
[75] S. Oviatt, "Multimodal interfaces," *Handbook of Human–Computer Interaction*, J. Jacko and A. Sears, Eds. Hillsdale, NJ: Lawrence Erlbaum Associates, 2002, pp. 286–304.
[76] M. Pantic, A. Pentland, A. Nijholt, and T. Huang, "Human computing and machine understanding of human behavior: A survey," in *Proc. 8th Int. Conf. Multimodal Interfaces (ICMI'06)*, Banff, Alberta, Canada, pp. 239–248.
[77] J. Pasquero. (2006). Survey on communication through touch, TR-CIM 06.04. Center for Intelligent Machines, McGill University, Montreal, Canada. [Online]. Available: http://www.cim.mcgill.ca/~haptic/pub/JP-CIM-TR-06.pdf
[78] J. Pasquero, J. Luk, S. Little, and K. Maclean, "Perceptual analysis of haptic icons: An investigation into the validity of cluster sorted MDS," in *Proc. 14th Symp. Haptic Interfaces for Virtual Environments and Teleoperator Systems (IEEE-VR'06)*, Alexandria, VA, pp. 437–444.
[79] J. Payette, V. Hayward, C. Ramstein, and D. Bergeron, "Evaluation of a force-feedback (haptic) computer pointing device in zero gravity," in *Proc. 5th ASME Symp. Haptic Interfaces for Virtual Environments and Teleoperated Systems*, 1996, vol. 58, pp. 547–553.

[80] R. W. Picard, *Affective Computing*. Cambridge, MA: MIT Press, 1997.
[81] I. Poupyrev and S. Maruyama, "Tactile interfaces for small touch screens," in *Proc. 16th Annu. ACM Symp. User Interface Software and Technology (UIST'03)*, pp. 217–220.
[82] L. Rosenberg, "Virtual fixtures: Perceptual overlays enhance operator performance in telepresence tasks," Ph.D. thesis, Stanford Univ., CA, 1994.
[83] J. Rovan and V. Hayward, "Typology of tactile sounds and their synthesis in gesture-driven computer music performance," *Trends in Gestural Control of Music*, M. Wanderley and M. Battier, Eds. Paris: IRCAM, 2000, pp. 297–320.
[84] S. Saga, N. Kawakami, and S. Tachi, "Learning effect of haptic teaching using opposite force presentation," *Nippon Kikai Gakkai Robotikusu, Mekatoronikusu Koenkai Koen Ronbunshu*, vol. 2005, p. 1P2-N-038, 2005.
[85] E. Sallnas, K. Rassmus-Grohn, and C. Sjostrom, "Supporting presence in collaborative environments by haptic force feedback," *Trans. Comput. Hum. Interact. (TOCHI)*, vol. 7, no. 4, pp. 461–476, 2000.
[86] N. B. Sarter, "Multimodal information presentation: Design guidance and research challenges," *Int. J. Ind. Ergonom.*, vol. 36, no. 5, pp. 439–445, 2006.
[87] R. M. Satava, "Emerging technologies for surgery in the 21st century," *Arch. Surg.*, vol. 134, no. 11, pp. 1197–1202, 1999.
[88] J. Smith and K. E. MacLean, "Communicating emotion through a haptic link: Design space and methodology," *Int. J. Hum. Comput. Stud. (IJHCS) (Special Issue on Affective Evaluation—Innovative Approaches)*, vol. 65, no. 4, pp. 376–387, 2007.
[89] S. S. Snibbe, K. E. MacLean, R. Shaw, J. B. Roderick, W. Verplank, and M. Scheeff, "Haptic metaphors for digital media," in *Proc. ACM Symp. User Interface Software and Technology (UIST'01)*, Orlando, FL pp. 199–208.
[90] C. Spence, M. E. R. Nicholls, N. Gillespie, and J. Driver, "Cross-modal links in exogenous covert spatial orienting between touch, audition, and vision," *Percept. Psychophys.*, vol. 60, no. 4, pp. 544–557, 1998.
[91] D. Stone, C. Jarrett, M. Woodroffe, and S. Minocha, *User Interface Design and Evaluation*. (*Interactive Technologies Series*). San Matio, CA: Morgan Kaufmann, 2005.
[92] R. Strong and B. Gaver, "Feather, scent, and shaker: Supporting simple intimacy," in *Proc. ACM Conf. Computer Supported Cooperative Work (CSCW'96)*, New York, NY pp. 29–30.
[93] C. Swindells, K. E. MacLean, K. S. Booth, and M. Meitner, "Exploring affective design for physical controls," in *Proc. ACM Conf. Human Factors in Computing Systems (CHI'07), CHI Lett.*, San Jose, CA, vol. 9, no. 1, pp. 933–942.
[94] C. Swindells, J. D. Smith, and K. E. MacLean, "An exploration of representations to aid design of haptic behaviours," in *Proc. CHI 2005 Workshop—Hands-on Haptics: Exploring Non-Visual Visualisation Using the Sense of Touch*, Portland, OR, pp. 5–8.
[95] H. Z. Tan, R. Gray, J. J. Young, and R. Traylor, "A haptic back display for attentional and directional cueing," *Haptics-e:Electron. J. Haptics Res.*, vol. 3, no. 1, 2003.
[96] H. Z. Tan and A. Pentland, "Tactual displays for sensory substitution and wearable computers," *Fundamentals of Wearable Computers and Augmented Reality,* W. Barfield and T. Caudell, Eds. Mahwah, NJ: Lawrence Erlbaum Associates, 2001, pp. 579–598.
[97] A. Tang, P. McLachlan, K. Lowe, R. S. Chalapati, and K. E. MacLean, "Perceiving ordinal data haptically under workload," in *Proc. 7th Int. Conf. Multimodal Interfaces (ICMI'05)*, Trento, Italy, pp. 244–251.
[98] R. Taylor, P. Jensen, L. Whitcomb, A. Barnes, R. Kumar, D. Stoianovici, P. Gupta, Z. Wang, E. Dejuan, and L. Kavoussi, "A steady-hand robotic system for microsurgical augmentation," *Int. J. Robotics Res.*, vol. 18, no. 12, pp. 1201–1210, 1999.
[99] C. L. Teo, E. Burdet, and H. P. Lim, "A robotic teacher of Chinese handwriting," in *Proc. 10th Symp. Haptic Interfaces for Virtual Environment and Teleoperator Systems, HAPTICS 2002*, pp. 335–341.
[100] D. Ternes, "Building large sets of haptic icons: Rhythm as a design parameter, and between-subjects MDS for evaluation," M.Sc. thesis, Univ. British Columbia, Vancouver, BC, Canada, 2007.
[101] B. Ullmer, H. Ishii, and D. Glass, "Mediablocks: Physical containers, transports, and controls for online media," in *Proc. Siggraph'98*, pp. 379–386.
[102] J. B. van Erp, H. A. Veen, C. Jansen, and T. Dobbins, "Waypoint navigation with a vibrotactile waist belt," *ACM Trans. Appl. Percept.*, vol. 2, no. 2, pp. 106–117, 2005.
[103] J. B. F. van Erp and M. M. A. Spapé, "Distilling the underlying dimensions of tactile melodies," in *Proc. Eurohaptics*, Dublin, Ireland, 2003, pp. 111–120.
[104] R. T. Verillo, "Vibration sensation in humans," *Music Percept.*, vol. 9, no. 3, pp. 281–302, 1992.
[105] W.-C. Wu, C. Basdogan, and M. A. Srinivasan, "Visual, haptic and bimodal perception of size and stiffness in virtual environments," in *Proc. 8th Annu. Symp. Haptic Interfaces for Virtual Environment and Teleoperator Systems, ASME/IMECE*, Nashville, TN, 1999, vol. 67, pp. 19–26.
[106] Y. Yanagida, M. Kakita, R. W. Lindeman, Y. Kume, and N. Tetsutani, "Vibrotactile letter reading using a low-resolution tactor array," in *Proc. 12th Int. Symp. Haptic Interfaces for Virtual Environment and Teleoperator Systems (HAPTICS'04)*, pp. 400–406.
[107] S. Yohanan, M. Chan, J. Hopkins, H. Sun, and K. E. MacLean, "Hapticat: Exploration of affective touch," in *Proc. 7th Int. Conf. Multimodal Interfaces (ICMI'05)*, Trento, Italy, pp. 244–251.

Karon E. MacLean received the B.S. degree in biological sciences and mechanical engineering from the Stanford University, California, in 1986 and the M.S. degree in mechanical engineering from MIT in 1988. She served as an engineer at the Center for Engineering Design, University of Utah, from 1988 to 1990. She received her Ph.D. in mechanical engineering from MIT in 1996. She was a member of the research staff of Interval Research Corp. from 1996 to 2000. She is now an associate professor at the Computer Science Department of the University of British Columbia. She is concerned with what haptic interfaces do; how they work; the way they feel, sound, and look; and how users will perceive them. She leads the Sensory Perception and Interaction Research Group (SPIN Lab) at the University of British Columbia.

Vincent Hayward received the Diplôme d'Ingénieur from École Centrale de Nantes in 1978. He received his Ph.D. in computer science from the University of Paris in 1981. He was Chargé de recherches at CNRS, France, from 1983 to 1986. Currently, he is professor of electrical and computer engineering at McGill University, Quebec, Canada. Hayward is interested in haptic interfaces design as well as applications, perception, and robotics. He is leading the Haptics Laboratory at McGill University. He is the recipient of the NASA Space Act Tech Brief Award (1991) and the E. (Ben) & Mary Hochhausen Award for Research in Adaptive Technology for Blind and Visually Impaired Persons (2002). He has several publications in the area of robotics and haptic interfaces, has served on several editorial boards, and has participated in the creation of spin-off companies.

Address for Correspondence: Karon MacLean, Department of Computer Science, The University of British Columbia, ICICS/CS Building, 201-2366 Main Mall, Vancouver, B.C., V6T 1Z4. E-mail: maclean@cs.ubc.ca.