Author's draft. The final version appeared in the IEEE Robotics and Automation Magazine, 15(1):104-119 (2008)

Do-It-Yourself Haptics, Part II: Interaction Design

Karon MacLean[a] and Vincent Hayward[b]

[a]SPIN Laboratory, The University of British Columbia, Vancouver, Canada

[b]Haptics Laboratory, McGill University, Montréal, Canada

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 which 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 which drive them, and specific examples which can be built "on the kitchen table". Here in Part II, we broach a topic which 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.

1 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 [72]. 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 simply not being 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;
- Providing 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 below.

1.1 Why Interaction Design Matters

There is not a computer user today without a collection of stories of user interfaces – generally graphical, as that is what we are surrounded with – that have annoyed, confused or stymied him or her. 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 are, rather, the intersection of bad user interface 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 user interface (UI) design is hard. It's not just a need for formal training and experience; this 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 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, for example, 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 (Human-Computer Interaction) design are all the more important when the modality is one that people are not accustomed to using in this way; and furthermore, one which 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.

1.2 Overview

In the remainder of this article, we'll start by considering the mapping between the cross-cutting roles it is thought that haptic feedback can serve in many different kinds of application spaces; and conversely, the human abilities and limitations which must be recognized, targeted or supported as these roles are developed (Section 2). In Section 3, we'll move on to a some design guidance which is especially relevant to haptic interactions, then close with a pair of case studies that illustrate contrasting approaches to actually doing it (Section 4).

2 Usable Roles for Haptic Feedback

Above, we listed some very broadly defined potential roles for haptic feedback. In a closer look, here we take a different cut; in each of 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 above. 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 [64] 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 non-technological 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 cellphone keyboards). We already see evidence of this here, e.g. in human acuity in abstract tactile message decoding. 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 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 short answer is that relevance of the source to *the task matters*, along with the source's trustworthiness [31].

2.1 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. tabletop) with non-grounded 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), so position sensing cannot be used to measure user intent; pressure might be more suitable.

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

2.2 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, 48, 57, 88]; see also Part I of this tutorial), we won't 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 mis-scaled 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 below).

An additional element is 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 [64]). In particular, it is necessary to understand a user's perceptual attributes in order 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 [57].

Finally, multimodal issues are almost always critical to attaining a realistic simulation result, in particular for renderings which 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, 45, 56, 106]).

2.3 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; this is discussed more below. *Guidance*, on the other hand, implies a more continuous engagement which is usually delivered through grounded force feedback for motor skills or, with lower resolution, via distributed tactors on the body for applications like 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 could be an expert system or another human; the latter is an instance of shared control or remote collaboration, which is also discussed more below. 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 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, 70] – 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 [32]

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 [52]; and requiring the student to *cancel* a reversed target force [85]. 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 [39, 54, 100]. These methods have not been directly compared with one another, 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 user and machine: an expert system has knowledge of the sensed and networked environment, databases, etc, 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 re-displayed 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 [49]; [99] illustrates the beginnings of force *sharing* during teleoperation).

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 [83]. Softening the guiding constraint turns this concept into mixed-initiative guidance: the user can choose to be guided, or "punch through" to do something else. Many variants of control-sharing using this concept have been tried (e.g. [35, 42, 53, 58, 74]; see [64] for a more thorough discussion). A sought-after first-order 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, there are however 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 [35]. 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 (e.g. [27, 41]). 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 forces in order to perform a physical task (as described in various contexts above) 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, 86], and conveyed the momentary degree of control balance between the partners [73]. In an explicitly social context, the *nature* of the force sharing impacts the sense of an interpersonal emotional connection [89].

2.4 Tactile Signaling in Multi-tasking Environments

Passive touch cues (which are presented to the observer's skin, rather than felt in response to active movements [37] can be used for notification of events and to create relatively unintrusive, 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 [80]). In this section, we'll therefore first mention issues relating to tactile design for multitasking, as well as typical methods and sites of delivery; and 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") which 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 email 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 [64].

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 vs. 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 (below).

Parameters used to encode content may also vary inherently in salience: e.g. in some schemes and for some display hardware, higher frequencies and/or amplitudes are perceived as "louder" than lower ones, and 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 which efficiently aid this task *at the same time* as optimizing meaning design.

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) [69, 77] and in modeling and detecting user mental/emotional state and interruptibility [33, 47].

Ambient Tactile Displays and the Human Body

Physical configuration and body site: It is generally 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 generally needed for more dextrous roles, the glabrous skin of the fingertips is rarely convenient as the delivery site; which for the most part leaves the less sensitive hairy skin [36]. Past examples, usually for simple signals, have used vests and belts [51, 96, 103]; back [96, 107]; and tongue [3].

Applications and contexts where hands can be used for background display include steering wheel [26], trackpoint [15], mouse [16, 17], and increasingly, mobile devices [55, 61, 82].

Active and passive touch: More fundamentally, [37] has argued that "passive touch … is atypical of normal tactile perception and that it leads the person to focus on the body surface;" whereas active touch is predominant in naturalistic environments where people are seeking information" [87]. Considering that convenient ambient tactile delivery sites are generally less sensitive skin *and* that 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, 91]. While signal complexity can be viewed as a continuum (defined either by information capacity in individual signals, or by 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 (a) 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), (b) only limited information need be conveyed, or (c) 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, non-glabrous 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 [59]. A recent review of this application space can be found in [97], backed by reviews of relevant aspects of tactile psychophysics [36, 50, 78]. Abstract tactile information transmission has centered on *haptic icons* or their equivalent: brief informative haptic signals (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 be very similar, others quite different, with no logical pattern); and the fact that, as we have observed in our early work here, 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 (e.g. [62], and see below).

One approach to increasing the controllability of the representational approach is to ascertain a set of basic primitives using careful testing, with the goal of then using them across contexts in a variety of situations (e.g. [98]. Another is for designers to carefully create their own codes, drawing on an existing knowledge base accessed by users [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, 34] 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. [64] summarizes psychophysical evidence for tactile acuity with respect to this kind of information transmission. There is some neural evidence of brain plasticity for users asked to pick up this skill after early childhood [36, 44].

What learning techniques will best exploit this plasticity? Taking encouragement from human ability to learn Braille after childhood [40] and guidance from how it is taught, we note that a first step is generally to develop the learner's tactual acuity. [5] list a 5-step process which 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. Immersion in rich and guided haptic experiences are 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 [46]. 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 [67].

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, 101, 104]. [64] summarizes the current status on dimensionality that has been found for various types of stimuli and display hardware.

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 minute session (depending on the set size), a few users can provide enough dissimilarity data about a stimulus set to create a map which reveals the dimensions along which the subjects perceive the stimuli relative to one another [62, 79, 101]. This map can be used to (a) 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, 62], and (b) 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. [68]), a multisensory reinforcement learning process is probably advantageous even to learn a stimulus which might later be invoked purely through the haptic channel.

In efforts to date, users have already demonstrated a good ability to learn associations which are either metaphorically matched by the designer (e.g. [13, 16, Tang, 2005 #426], deliberately arbitrary [29, 30] 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].

2.5 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 [62]. 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: the *product* of typing (the printed word) can be highly emotive in every way, both visually (ask a typesetter) and in linguistic meaning. But, 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. E.g. 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, 84, 105]. This is further linked to a consistency or closing of the control loop – a mechanical interaction between the subject and the sound source [38, 60]. However, computer-controlled grounded forces bring constraints: 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, both from the standpoint of the fundamental interactions themselves, and their mechatronic implementation [11, 20, 63, 75]. [64] links into the 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. A common feature is strong individuation of "instrument" to application, i.e. type of music to be created and gestures employed: these are not general-purpose devices.

2.6 Haptic Affect

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 which impacts stress and usability [71]. It also forms the basis of a new, sophisticated type of interface based on *affective computing* [81], where the computer senses and displays determine and elicit particular emotional experiences form 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 [64]. Here, we mention two potential roles for haptic design.

Design for Feel

Consider the direct affective response that *feels* produce 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 affect (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 [94]. Eventually, this line of research should result in heuristics that will guide, for example, choice of feel for a given control action. For now, a 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 by [93], 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" (review, [43]). 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 [64]). 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, 89] in at least simplified contexts; in other work we are building a touch sensing-and-display platform to study this in a less constrained environment [108].

3 Haptic Interaction Design Practices

There is a wealth of information on best practices for user interface design – textbooks [8, 24, 92], courses, conferences and journals. There is also a growing literature on principles for *multimodal* interface design, which is relevant here [76, 87]. But what is special about the process of designing *haptics* into interfaces? Or, even better, designing the interface itself around the idea of physical interaction?

3.1 Technocentric vs. 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're 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 certain mentality.

- Have Technology, Seeking Need: If you're looking for an application that will show off your device's special features, you (a) might have a fruitless search; or (b) 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 that cool gadget. It's important to watch and listen to people and notice where they struggle, and holding an open mind: perhaps your original idea isn't the right one, but the problem is real and understanding it will guide you to a different and better solution.

+ 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 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.

3.2 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 which 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 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, say, 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, so you 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. I believe the most informative compromise is often to compare best-of-breed versions and focus on collecting and analyzing rich observational data, in contrast to a hypothesis-test emphasizing quantitative performance measures and statistical differences.

Evaluation in the Middle of the Learning Curve: The playing field isn't always level. For example, with tactile signaling, 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- to 30- minute training period. It can be difficult to determine whether an innovation has intrinsic value, nor 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 [95].

3.3 Importance of Rapid Prototyping and Haptic Representations

A well-recognized principle of prototyping in *all* disciplines is to iterate at increasing levels of detail, whether creating a piece of software, 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 rigidified than in early, conceptual stages. For user interfaces, the truth of this maxim grows: while there are some trustworthy heuristics, it is difficult to predict user response to any kind of novelty – modeling, simulation or established rules are restricted. For *haptic* user interfaces, the unfamiliarity and the combination of hardware and software design further amplifies this.

Minimalism: Prototyping user interfaces 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 "finalized" to a user, who will be less likely to challenge or suggest modifications.

Modular Prototyping: The primary objective of a prototype is to *get your design question (starting with big ones and proceeding to more detailed ones) answered with as little effort as possible, then 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 non-actuated 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 the 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 and risky, but presumably 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 below.

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. When possible, advance two or three different paths which 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, 92]).

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 handheld devices which cannot yet be built with sufficient compactness and power efficiency to actually be portable. But 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 found a way to use it effectively.

3.4 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. But there are a couple of 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 pre-digital 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 below, 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 (with similarities to one which users might have once used in the real world) [63, 90]

Navigating Modes

Haptic feedback is often proposed as a solution for *modal interfaces:* those where the interface can be in different states, and a command 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 LCD display) has possibilities here because unlike a physical knob, it can be reprogrammed appropriately for the current mode, just like the graphical display. But, 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 mode. Our current hardware state-of-art (point-based interaction for force feedback, at least) means that usually you have to explore an environment serially, e.g. turn the knob 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 which 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 receiving (passively) ongoing physical feedback about state which does not require continual system interrogation. Physical metaphor is a good way to enter into this idea, because it is how real handheld 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 this case study allows one to move seamlessly between these states, and the "tick-tick" 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. And if the set is represented in the right way, you can skim over the discrete items themselves like a texture, feeling for the variation which indicates the item property you're looking for.

4 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 variety of principled design mechanisms. For authenticity and detail as well as brevity and focus, they are chosen from the authors' own experience.

4.1 Force Feedback Knob: Continuous and Discrete Handheld Media Control

Along with digitization of once-tangible tasks and (with ubiquitous computing) *controllers* 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 which relied on *metaphor* to create generalizable but experience-grounded interactions for a hand-held media controller [65, 66, 90][1], beginning with some relevant principles and observations. This case also illustrates the "modular prototyping" principle described in Section 3. Starting from the ideas of metaphor-based design and discrete/continuous media modes, we set out to build a handheld "universal home media controller" which would leverage the apparent utility of modal interaction for various types of media in a consistent way, while making state transparently clear.

Inspiring Metaphor: We tried a lot of metaphors! And we ended up using several. For example, one which users found compelling (it felt good and aided their navigation) was a virtual "clutch" through which the user interacted with a heavy reel of "film" which 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 of stop/start buttons.

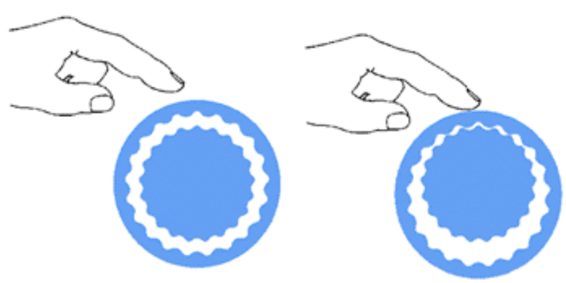

Figure 1: Virtual Clutch metaphor for the force feedback media controller. The knob is equipped with a crude pressure sensor. When the user presses down on the actual knob (which is associated with the outer wheel in this figure) the heavy inner wheel (virtual) is engaged and can be spun up. When the 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 which fuse into a texture as the speed increases.

[1] The first stages of this project were performed at Interval Research Corporation (Palo Alto, CA) during 1998-1999 by a design team lead by the first author. Later stages were conducted as student projects at UBC, also directed by the first author.

Technical grounding: A technical path for this was suggested by "tangible interfaces", where "tagged" arbitrary objects (e.g. using Radio Frequency Identification, or RFID) can be used to issue commands to a computer [102]. Observing that tagged objects are well suited for issuing digital commands but not for exerting continuous control, we combined the two through the use of "Tagged Handles" [66, MacLean, 2002 #267].

Prototype-Driven Design Steps

Figure 2 illustrates several successive prototypes in an iterative conceptual and engineering evolution: in the 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 prototype (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, or selection of radio versus recorded content, volume versus navigational control. No attempt was made at usability, e.g. the handles did not suggest their function.

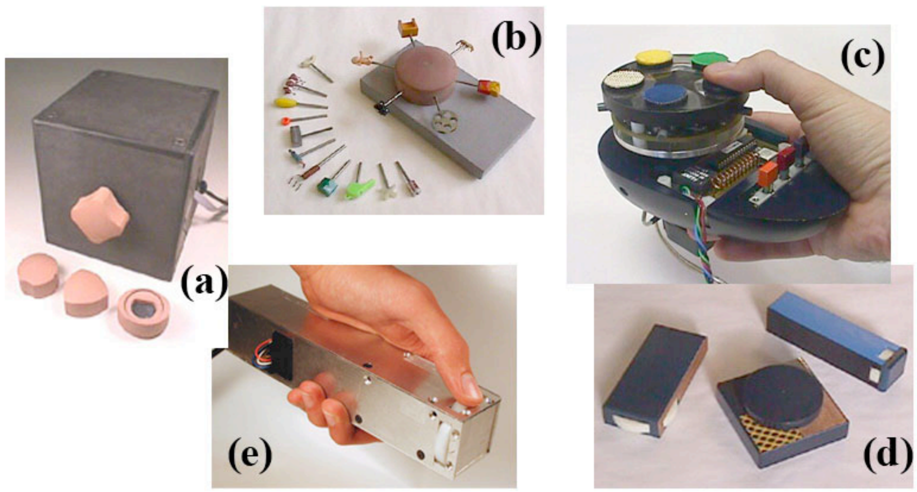

Figure 2: Representative haptic media controller design iterations. Clockwise from upper left, we see (a) the initial "Tagged Handles engineering concept prototype, (b) a representative conceptual prototype, (c) a later technical prototype (oversized), (d) a set of non-functional concept prototypes which address the problems of (c) and (e) another engineering mockup

Prototype (b) shows one example of many ideas explored at a mostly conceptual level in Figure 3: a non-functioning prototype showing one way that discrete handles (inspired by a charm bracelet) could informatively indicate their function solve the practical problem of getting lost – the handle is selected from a wheel instead of 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 3: Early prototype for the handheld media controller project. Found objects and state-of-art examples, Lego + rubber band transmissions, whimsical and serious non-functioning concepts, and narrowly targeted functional engineered prototypes.

Prototype (c) shows 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 handheld 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 (d) inverted this idea: a four-sided object with texturally marked sides and an active thumbwheel knew which face was "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, from another select channel, etc. Finally, (e) is another engineering prototype of this final idea [65].

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 (largely 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 up-front observation, and the next step would have been a usability study[2].

4.2 Vibrotactile Background Signals: Mediation of Turntaking in Remote Collaboration

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 above (Section 2.4), and culminating in an *observational user study*. These steps are more fully described in [16, 17] and we summarize some key points here.

Application: When non-co-located, collaborating users wish to jointly modify a shared object displayed on their local screens – whether a text document, a CAD 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 co-located situations[3].

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

---

[2] … if the host company hadn't vaporized in the 2000 tech bust.

[3] Our own guess is that even the usual non-verbal cues available in co-located 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?

methods could support – because the latter were also being used in the collaborative task. We further wondered if the ability to make 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 ear-shot 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 you 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.

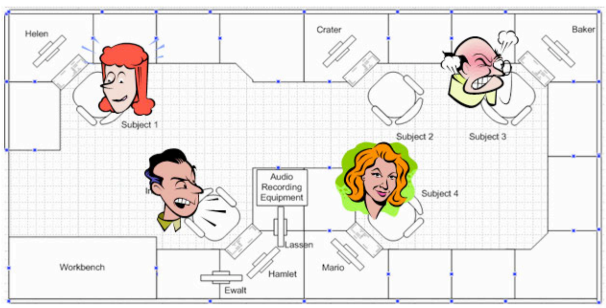

Figure 4: Experiment setup for observational study of the turntaking protocol. The four group members were placed out of direct eye-shot and wore noise-canceling headphones; all vocal communications occurred through a sound system.

Figure 5: Vibrotactile mouse used to display the haptic icons used in the turntaking protocol. Two buttons were added to the side, to enable special protocol features; such buttons were available in other mice at the time, but not the vibrotactile one.

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; the initial set was a bit different. We used metaphor-based design on the assumption 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 inserting or removing a hardware device from your computer.

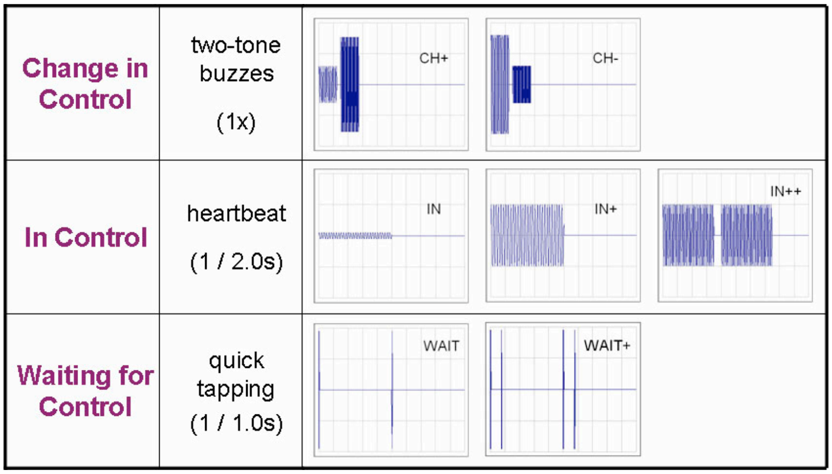

Figure 6: Final set of haptic icons used in the turntaking protocol.

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

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

Step II: We *perceptually adjusted* the icon set using the MDS technique mentioned above, 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.

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 if the urgency-based protocol. Following this test, we adjusted some of the signals a bit more to get the desired salience patterns. Subjects learned the 7 mediating icons easily in 3 minutes, and maintained 97% accuracy of identification under substantial multimodal workload.

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

Step IV: Finally, we mounted the group observational study, and learned lots (read the paper). Through a combination of performance and subjective measures we did, for example, confirm that the haptic signals were utilized in a graded (i.e. appropriate), 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 (our visual implementation allowed the queue to be minimized; it provided requestor identity when opened, which the tactile cues did not).

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. But 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 "rich" – i.e. 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.

5 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 follow, and pitfalls to avoid, 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 Part I's electromechanical design principles, will lower the "barrier to entry" for this exciting young field, and foment many new ideas - usable ones!

6 References

[1] Adams, R. J., Klowden, D., and Hannaford, B., "Virtual Training for a Manual Assembly Task," *haptics-e*, vol. 2:2, 2001.

[2] Avizzano, C. A., Solis, J., and Bergamasco, M., "Teaching to Write Japanese Characters Using a Haptic Interface," in Proc. of the 10th Symposium on Haptic interfaces For Virtual Environment and Teleoperator Systems, Washington, DC, 2002.

[3] Bach-y-Rita, P., Kaczmarek, K. A., Tyler, M. E., and Garcia-Lara, M., "Form Perception with a

49-Point Electrotactile Stimulus Array on the Tongue: A Technical Note," *J. Rehab. Res. Dev*, vol. 35, pp. 427-430, 1998.

[4] Bailenson, J. N., Yee, N., Brave, S., Merget, D., and Koslow, D., "Virtual Interpersonal Touch: Expressing and Recognizing Emotions through Haptic Devices," *Human-Computer Interaction*, 2007, in press.

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

[6] Basdogan, C., Ho, C. H., Srinivasan, M., and Slater, M., "An Experimental Study on the Role of Touch in Shared Virtual Environments," *ACM Transactions on Computer-Human Interaction (TOCHI): Special issue on human-computer interaction and collaborative virtual environments*, vol. 7:4, pp. 443–460, 2000.

[7] Beamish, T., MacLean, K., and Fels, S., "Manipulating Music: Multimodal Interaction for Djs," in Proc. of ACM Conf on Human Factors in Computing Systems (CHI '04), CHI Letters, vol 6, no. 1, Vienna, Austria, pp. 327 - 334, 2004.

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

[9] Bertelson, P. and de Gelder, B., "The Psychology of Crossmodal Attention," in *Crossmodal Space and Crossmodal Attention*, C. Spence and J. Driver, Eds.: Oxford University Press Inc, 2004, pp. 141-177.

[10] Blake, S. J., Teaching Reading in Braille, 2007.

http://blindness.growingstrong.org/ed/aa010401a.htm.

[11] Bongers, A. J., "Interactivation: Towards an E-Cology of People, Our Technological Environment, and the Arts," Proefschrift Amsterdam, VU., Amsterdam 2006-12, 2006.

[12] Brave, S. and Dahley, A., "Intouch: A Medium for Haptic Interpersonal Communication," in Proc. of the Conf. on Human Factors in Computing Systems (CHI'97), pp. 363-364, 1997.

[13] Brown, L. M., Brewster, S. A., and Purchase, H. C., "Multidimensional Tactons for Non-Visual Information Presentation in Mobile Devices," in Proc. of the 8th conference on Human-computer interaction with mobile devices and services, Helsinki, Finland, 2006

[14] Camarillo, D. B., Krummel, T. M., and Salisbury, J. K., "Robotic Technology in Surgery: Past, Present, and Future," *American Journal of Surgery*, vol. 188:4A, pp. 2S-15S, 2004.

[15] Campbell, C. S., Zhai, S., May, K. W., and Maglo, P., "What You Feel Must Be What You See: Adding Tactile Feedback to the Trackpoint," in Proc. of Interact'99, Edinburgh, UK, pp. 383-390, 1999.

[16] Chan, A., MacLean, K. E., and McGrenere, J., "Learning and Identifying Haptic Icons under Workload," in Proc. of First Joint Eurohaptics Conf. and Symp. on Haptic Interfaces for Virtual Environment and Teleoperator Systems (WHC 2005, IEEE-VR2005), Pisa, Italy, pp. 432-439, 2005.

[17] Chan, A., MacLean, K. E., and McGrenere, J., "Designing Haptic Icons to Support Collaborative Turn-Taking," *International Journal of Human Computer Studies*, to appear.

[18] Chang, A., O'Modhrain, S., Jacob, R., Gunther, E., and Ishii, H., "Comtouch: Design of a Vibrotactile Communication Device," in Proc. of Conf. on Designing Interactive Systems (DIS'02), London, pp. 312-320, 2002.

[19] Choi, S. and Tan, H. Z., "An Analysis of Perceptual Instability During Haptic Texture Rendering," in Proc. of 10th Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems, Los Alamitos, CA, 2002.

[20] Cook, P., "Principles for Designing Computer Music Controllers," in Proc. of Conference on New Interfaces for Musical Expression, Seattle, Washington, 2001.

[21] Craig, J. C., "Vibrotactile Pattern Perception: Extraordinary Observers," *Science*, vol. 196:4288, pp. 450-452, 1977.

[22] Delp, S., Loan, P., Basdogan, C., and Rosen, J. M., "Surgical Simulation: An Emerging Technology for Emergency Medical Training," *Presence: Teleoperators and Virtual Environments*, vol. 6, pp. 147-159, 1997.

[23] DiFilippo, D. and Pai, D. K., "Contact Interaction with Integrated Audio and Haptics," in Proc. of Int'l Conf on Auditory Display (ICAD'00), 2000.

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

[25] Dixon, S. R., Wickens, C. D., and McCarley, J. S., "On the Independence of Compliance and Reliance: Are Automation False Alarms Worse Than Misses?," *Human Factors*, vol. 49:4, pp. 564-572, 2007.

[26] Enriquez, M., Afonin, O, Yager, B, MacLean, K, "A Pneumatic Tactile Notification System for the Driving Environment," in Proc. of Workshop on Perceptive User Interfaces (PUI '01), Orlando, FL, pp. 1-7, 2001.

[27] Enriquez, M. and MacLean, K., "Impact of Haptic Warning Signal Reliability in a Time-and-Safety-Critical Task," in Proc. of the 12th Annual Symposium on Haptic Interfaces for Virtual Environments and Teleoperator Systems, IEEE-VR2004, Chicago, pp. 407-415, 2004.

[28] Enriquez, M. and MacLean, K. E., "Persistent Recall of Haptic Icons Created with Arbitrary Versus User-Selected Associatons," *submitted to the IEEE Symposium on Haptic Interfaces for Virtual Environments and Teleoperator Systems (HAPTICS 2008)*, 2008.

[29] Enriquez, M. and MacLean, K. E., "The Role of Choice in Longitudinal Recall of Meaningful Tactile Signals," in Proc. of submitted to IEEE Symposium on Haptic Interfaces (HAPTICS-08), 2008.

[30] Enriquez, M., MacLean, K. E., and Chita, C., "Haptic Phonemes: Basic Building Blocks of Haptic Communication," in Proc. of 8th International Conference on Multimodal Interfaces (ICMI 2006), Banff, Canada, 2006.

[31] Ernst, M. O. and Banks, M. S., "Humans Integrate Visual and Haptic Information in a Statistically Optimal Fashion," *Nature*, vol. 415, pp. 429–433, 2002.

[32] Feygin, D., Keehner, M., and Tendick, F., "Haptic Guidance: Experimental Evaluation of a Haptic Training Method for a Perceptual Motor Skill," in Proc. of 10th Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems, Los Alamitos, CA, 2002.

[33] Fogarty, J., Hudson, S. E., Atkeson, C. G., Avrahami, D., Forlizzi, J., Kiesler, S., Lee, J. C., and Yang, J., "Predicting Human Interruptibility with Sensors," *ACM Transactions on Computer-Human Interaction*, vol. 12:1, pp. 119-146, 2005.

[34] Fogg, B., Cutler, L. D., Arnold, P., and Eisbach, C., "Handjive: A Device for Interpersonal Haptic Entertainment," in Proc. of Conf. on Human Factors in Computing Systems, pp. 57-64, 1998.

[35] Forsyth, B. and MacLean, K. E., "Predictive Haptic Guidance: Intelligent User Assistance for the Control of Dynamic Tasks," *IEEE Transactions on Visualization and Computer Graphics*, vol. 12:1, pp. 103-113, 2006.

[36] Gallace, A., Tan, H. Z., and Spence, C., "The Body Surface as a Communication System: The State of the Art after 50 Years," *Presence: Teleoperators & Virtual Environments*, 2007.

[37] Gibson, J. J., "Observations on Active Touch," *Psychological Review*, vol. 69, pp. 477-490, 1962.

[38] Gillespie, B., "The Touchback Keyboard," in Proc. of the International Computer Music Conference, San Jose, CA, pp. 447--448, 1992.

[39] Gillespie, B., O'Modhrain, S., P. Tang, C. P., and Zaretsky, D., "The Virtual Teacher," in Proc. of Proceedings of the ASME Dynamic Systems and Control Division, DSC-Vol. 64, pp. 171-178, 1998.

[40] Grant, A., Thiagarajah, M., and Sathian, K., “Tactile Perception in Blind Braille Readers: A Psychophysical Study of Acuity and Hyperacuity Using Gratings and Dot Patterns.,” *Perceptual Psychophysics*, vol. 62:2, pp. 301-12, 2000.

[41] Griffin, W. B., Provancher, W. R., and Cutkosky, M. R., “Feedback Strategies for Shared Control in Dexterous Telemanipulation,” in Proc. of Intelligent Robots and Systems, 2003. (IROS 2003). Proceedings. 2003 IEEE/RSJ International Conference on, 3, pp. 2791-2796 vol.3, 2003.

[42] Griffiths, P. and Gillespie, R. B., “Shared Control between Human and Machine: Haptic Display of Automation During Manual Control of Vehicle Heading,” in Proc. of the 12th International Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems (HAPTICS'04) pp. 358-366, 2004.

[43] Haans, A. and IJsselsteijn, W., “Mediated Social Touch: A Review of Current Research and Future Directions,” *Virtual Reality*, vol. 9, pp. 149-159, 2006.

[44] Hamilton, R. H. and Pascual-Leone, A., “Cortical Plasticity Associated with Braille Learning,” *Trends in Cognitive Sciences*, vol. 2, pp. 168-174, 1998.

[45] Hayward, V., “A Brief Taxonomy of Tactile Illusions and Demonstrations That Can Be Done in a Hardware Store,” *Brain Research Bulletin (special issue on Robotics and Neuroscience)*, 2007, in press.

[46] Hollins, M., Bensmaïa, S., Karlof, K., and Young, F., “Individual Differences in Perceptual Space for Tactile Textures: Evidence from Multidimensional Scaling,” *Perception & Psychophysics*, vol. 62, pp. 1534-1544, 2000.

[47] Horvitz, E., Kadie, C., Paek, T., and Hovel, D., “Models of Attention in Computing and Communication: From Principles to Applications,” *Communications of the ACM*, vol. 46:3, pp. 52-59, 2003.

[48] Howe, R. D. and Matsuoka, Y., “Robotics for Surgery,” *Annual Review of Biomedical Engineering*, vol. 1:211-240, 1999.

[49] Jacobsen, S. C., Smith, F. M., Backman, D. K., and Iversen, E. K., “High Performance, High Dexterity, Force Reflective Teloperator Ii,” in Proc. of ANS Topical Meeting on Robotics and Remote Systems, Albuquerque, NM, 1991.

[50] Jones, L. and Sarter, N., “Tactile Displays: Guidance for Their Design and Application,” *Human Factors: The Journal of Human Factors and Ergonomic Society*, vol. to appear, 2007.

[51] Jones, L. A., Nakamura, M., and Lockyer, B., “Development of a Tactile Vest,” in Proc. of Haptic Interfaces for Virtual Environment and Teleoperator Systems (HAPTICS '04), Chicago, pp. 82-89, 2004.

[52] Kikuuwe, R. and Yoshikawa, T., “Haptic Display Device with Fingertip Presser for Motion/Force Teaching to Human,” in Proc. of IEEE Int'l Conf. Robotics and Automation, pp. 868–873, 2001.

[53] Kragic, D., Marayong, P., Li, M., Okamura, A. M., and Hager, G. D., “Human Machine Collaborative Systems for Microsurgical Applications,” *International Journal of Robotics Research*, vol. 24:9, pp. 731-742, 2005.

[54] Krebs, H. I., Hogan, N., Aisen, M. L., and Volpe, B. T., “Robot-Aided Neuro-Rehabilitation,” *IEEE - Transactions on Rehabilitation Engineering*, vol. 6:1, pp. 75-87, 1998.

[55] Leung, R., MacLean, K. E., Bertelsen, M. B., and Saubhasik, M., “Evaluation of Haptically Augmented Touchscreen Gui Elements under Cognitive Load,” in Proc. of 9th International Conference on Multimodal Interfaces (ICMI '07), 2007.

[56] Levitin, D. J., MacLean, K. E., Mathews, M. V., Chu, L., and Jensen, E. R., “The Perception of Cross-Modal Simultaneity,” *Computing Anticipatory Systems*, vol. 517, pp. 323-329, 2000.

[57] Lin, M. and Otaduy, M., “Haptic Rendering: Foundations, Algorithms and Applications,” A K Peters, Ltd, in press.

[58] Lin, M. and Taylor, R. H., “Spatial Motion Constraints in Medical Robot Using Virtual Fixtures Generated by Anatomy,” in Proc. of IEEE Int'l Conf on Robotics and Automation (ICRA '04), 2, pp. 1270-1275, 2004.

[59] Linvill, J. G. and Bliss, J. C., “A Direct Translation Reading Aid for the Blind,” in Proc. of Institute of Electrical and Electronics Engineers, 54, pp. 40--51, 1966.

[60] Luciani, A., Florens, J., and Castagne, N., “From Action to Sound: A Challenging Perspective for Haptics,” in Proc. of First Joint Eurohaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems (WHC'05), pp. 592-595, 2005.

[61] Luk, J., Pasquero, J., Little, S., MacLean, K., Levesque, V., and Hayward, V., "A Role for Haptics in Mobile Interaction: Initial Design Using a Handheld Tactile Display Prototype," in Proc. of ACM Conference on Human Factors in Computing Systems (CHI '06), CHI Letters, vol. 8, no. 1, Montreal, Canada, pp. 171-180, 2006.

[62] MacLean, K. and Enriquez, M., "Perceptual Design of Haptic Icons," in Proc. of Proceedings of Eurohaptics, Dublin, Ireland, 2003.

[63] MacLean, K. E., "Designing with Haptic Feedback," in Proc. of IEEE Robotics and Automation (ICRA '00), San Francisco, CA, 1, pp. 783-788, 2000.

[64] MacLean, K. E., "Haptics in the Wild: Interaction Design for Everyday Interfaces," in *To Appear in Reviews of Human Factors and Ergonomics*, M. Carswell, Ed. Santa Monica, CA: Human Factors and Ergonomics Society, 2008.

[65] MacLean, K. E., Shaver, M. J., and Pai, D. K., "Handheld Haptics: A Usb Media Controller with Force Sensing," in Proc. of the IEEE VR2002 10th Symp. on Haptic Interfaces for Virtual Environment and Teleoperator Systems, Orlando, FL, pp. 311-318, 2002.

[66] MacLean, K. E., Snibbe, S. S., and Levin, G., "Tagged Handles: Merging Discrete and Continuous Control," in Proc. of ACM Conference on Human Factors in Computing Systems (CHI '00), CHI Letters, vol. 2, no. 1, The Hague, Netherlands, pp. 225-232, 2000.

[67] Millar, S., *Reading by Touch*: Routledge, 1997.

[68] Millar, S., "Memory in Touch," *Psicothema*, vol. 11, pp. 747-767, 1999.

[69] Moran, T. and Dourish, P., "Introduction to This Speical Issue on Context-Aware Computing," *Human-Computer Interaction*, vol. 16:2, pp. 87--95, 2001.

[70] Morris, D., Tan, H., Barbagli, F., Chang, T., and Salisbury, K., "Haptic Feedback Enhances Force Skill Learning," in Proc. of Proceedings of IEEE World Haptics, 2007.

[71] Norman, D. A., *Emotional Design: Why We Love (or Hate) Everyday Things*. New York: Basic Books, 2003.

[72] Norman, D. A., "The Next Ui Breakthrough, Part 2: Physicality," *ACM Interactions*, vol. 14:3, pp. 44-45, 2007.

[73] Nudehi, S. S., Mukherjee, R., and Ghodoussi, M., "A Shared-Control Approach to Haptic Interface Design for Minimally Invasive Telesurgical Training," *Control Systems Technology, IEEE Transactions on*, vol. 13:4, pp. 588-592, 2005.

[74] O'Malley, M. K., Gupta, A., Gen, M., and Li, Y., "Shared Control in Haptic Systems for Performance Enhancement and Training," *Journal of Dynamic Systems, Measurement, and Control*, vol. 128:1, pp. 75-85, 2006.

[75] O'Sullivan, D. and Igoe, T., *Physical Computing*. Boston, USA: Thomson Course Technology, 2004.

[76] Oviatt, S., "Multimodal Interfaces," in *Handbook of Human–Computer Interaction*, J. Jacko and A. Sears, Eds. Hillsdale, NJ: Lawrence Erlbaum Associates, 2002

[77] Pantic, M., Pentland, A., Nijholt, A., and Huang, T., "Human Computing and Machine Understanding of Human Behavior: A Survey," in Proc. of the 8th international Conference on Multimodal interfaces (ICMI '06), Banff, Alberta, Canada, pp. 239-248, 2006.

[78] Pasquero, J., "Survey on Communication through Touch," Center for Intelligent Machines, McGill University, Montreal, Canada TR-CIM 06.04, 2006. http://www.cim.mcgill.ca/~haptic/pub/JP-CIM-TR-06.pdf.

[79] Pasquero, J., Luk, J., Little, S., and Maclean, K., "Perceptual Analysis of Haptic Icons: An Investigation into the Validity of Cluster Sorted Mds," in Proc. of 14th Symposium on Haptic Interfaces for Virtual Environments and Teleoperator Systems (IEEE-VR'06), Alexandria, Virginia, 2006.

[80] Payette, J., Hayward, V., Ramstein, C., and Bergeron, D., "Evaluation of a Force-Feedback (Haptic) Computer Pointing Device in Zero Gravity," in Proc. of Fifth ASME Symp. on Haptic Interfaces for Virtual Environments and Teleoperated Systems, DSC-Vol. 58, pp. 547-553, 1996.

[81] Picard, R. W., "Affective Computing." Cambridge, MA: MIT Press, 1997.

[82] Poupyrev, I. and Maruyama, S., "Tactile Interfaces for Small Touch Screens," in Proc. of 16th annual ACM Symposium on User Interface Software and Technology (UIST'03), pp. 217–220, 2003.

[83] Rosenberg, L., *Virtual Fixtures: Perceptual Overlays Enhance Operator Performance in Telepresence Tasks,* Ph.D. Thesis, Stanford University, 1994.

[84] Rovan, J. and Hayward, V., “Typology of Tactile Sounds and Their Synthesis in Gesture-Driven Computer Music Performance,” in *Trends in Gestural Control of Music*, M. Wanderley and M. Battier, Eds. Paris: IRCAM, 2000, pp. 297--320.

[85] Saga, S., Kawakami, N., and Tachi, S., “Learning Effect of Haptic Teaching Using Opposite Force Presentation,” *Nippon Kikai Gakkai Robotikusu, Mekatoronikusu Koenkai Koen Ronbunshu*, vol. 2005, pp. 1P2-N-038, 2005.

[86] Sallnas, E., Rassmus-Grohn, K., and Sjostrom, C., “Supporting Presence in Collaborative Environments by Haptic Force Feedback,” *Transactions on Computer-Human Interaction (TOCHI)*, vol. 7:4, pp. 461-476, 2000.

[87] Sarter, N. B., “Multimodal Information Presentation: Design Guidance and Research Challenges,” *International Journal of Industrial Ergonomics*, vol. 36:5, pp. 439-445, 2006.

[88] Satava, R. M., “Emerging Technologies for Surgery in the 21st Century,” *Arch Surg*, vol. 134:11, pp. 1197-1202, 1999.

[89] Smith, J. and MacLean, K. E., “Communicating Emotion through a Haptic Link: Design Space and Methodology,” *International Journal of Human Computer Studies (IJHCS), Special issue on Affective Evaluation--Innovative Approaches*, vol. 65:4, pp. 376-387, 2007.

[90] Snibbe, S. S., MacLean, K. E., Shaw, R., Roderick, J. B., Verplank, W., and Scheeff, M., “Haptic Metaphors for Digital Media,” in Proc. of ACM Symp. on User Interface Software & Technology (UIST '01), Orlando, FL, pp. 199-208, 2001.

[91] Spence, C., Nicholls, M. E. R., Gillespie, N., and Driver, J., “Cross-Modal Links in Exogenous Covert Spatial Orienting between Touch, Audition, and Vision,” *Perception & Psychophysics*, vol. 60, pp. 544-557, 1998.

[92] Stone, D., Jarrett, C., Woodroffe, M., and Minocha, S., *User Interface Design & Evaluation*: Elsevier, Inc., 2005.

[93] Strong, R. and Gaver, B., “Feather, Scent, and Shaker: Supporting Simple Intimacy,” in Proc. of ACM Conference on Computer Supported Cooperative Work (CSCW '96) (short paper), New York, NY, pp. 29-30, 1996.

[94] Swindells, C., MacLean, K. E., Booth, K. S., and Meitner, M., “Exploring Affective Design for Physical Controls,” in Proc. of ACM Conference on Human Factors in Computing Systems (CHI '07), CHI Letters, vol. 9, no. 1, San Jose, CA, pp. 933-942, 2007.

[95] Swindells, C., Smith, J. D., and MacLean, K. E., “An Exploration of Representations to Aid Design of Haptic Behaviours,” in Proc. of CHI 2005 Workshop - Hands-on Haptics: Exploring Non-Visual Visualisation Using the Sense of Touch, Portland, OR, pp. 5-8, 2005.

[96] Tan, H. Z., Gray, R., Young, J. J., and Traylor, R., “A Haptic Back Display for Attentional and Directional Cueing,” *Haptics-e: The Electronic Journal of Haptics Research*, vol. 3:1, pp. 20 pp., 2003.

[97] Tan, H. Z. and Pentland, A., “Tactual Displays for Sensory Substitution and Wearable Computers,” in *Fundamentals of Wearable Computers and Augmented Reality*, W. Barfield and T. Caudell, Eds. Mahwah, NJ: Lawrence Erlbaum Associates, 2001.

[98] Tang, A., McLachlan, P., Lowe, K., Chalapati, R. S., and MacLean, K. E., “Perceiving Ordinal Data Haptically under Workload,” in Proc. of 7th International Conference on Multimodal Interfaces (ICMI '05), Trento, Italy, pp. pages 244-251, 2005.

[99] Taylor, R., Jensen, P., Whitcomb, L., Barnes, A., Kumar, R., Stoianovici, D., Gupta, P., Wang, Z., Dejuan, E., and Kavoussi, L., “A Steady-Hand Robotic System for Microsurgical Augmentation,” *The International Journal of Robotics Research*, vol. 18:12, pp. 1201-1210, 1999.

[100] Teo, C. L., Burdet, E., and Lim, H. P., “A Robotic Teacher of Chinese Handwriting,” in Proc. of Haptic Interfaces for Virtual Environment and Teleoperator Systems, 2002. HAPTICS 2002. Proceedings. 10th Symposium on, pp. 335-341, 2002.

[101] Ternes, D., *Building Large Sets of Haptic Icons: Rhythm as a Design Parameter, and between-Subjects Mds for Evaluation,* M.Sc. Thesis, University of British Columbia, 2007.

[102] Ullmer, B., Ishii, H., and Glass, D., “Mediablocks: Physical Containers, Transports, and Controls for Online Media,” in Proc. of Siggraph '98, 1998.

[103] van Erp, J. B., Veen, H. A., Jansen, C., and Dobbins, T., “Waypoint Navigation with a Vibrotactile Waist Belt,” *ACM Transactions on Applied Perception*, vol. 2:2, pp. 106-117, 2005.

[104] van Erp, J. B. F. and Spapé, M. M. A., "Distilling the Underlying Dimensions of Tactile Melodies," in Proc. of Eurohaptics, Dublin, Ireland, pp. 111-120, 2003.

[105] Verillo, R. T., "Vibration Sensation in Humans," *Music Perception*, vol. 9:3, pp. 281-302, 1992.

[106] Wu, W.-C., Basdogan, C., and Srinivasan, M. A., "Visual, Haptic and Bimodal Perception of Size and Stiffness in Virtual Environments," in Proc. of the 8th Ann. Symp. on Haptic Interfaces for Virtual Environment and Teleoperator Systems, ASME/IMECE, Nashville, TN, DSC:67, pp. 19-26, 1999.

[107] Yanagida, Y., Kakita, M., Lindeman, R. W., Kume, Y., and Tetsutani, N., "Vibrotactile Letter Reading Using a Low-Resolution Tactor Array," in Proc., 00, pp. 400, 2004.

[108] Yohanan, S., Chan, M., Hopkins, J., Sun, H., and MacLean, K. E., "Hapticat: Exploration of Affective Touch," in Proc. of 7th International Conference on Multimodal Interfaces (ICMI '05), Trento, Italy, pp. pages 244-251, 2005.