

Perceiving Ordinal Data Haptically Under Workload

Anthony Tang[†] Peter McLachlan[‡] Karen Lowe[‡] Chalapati Rao Saka[‡] Karon MacLean[‡]
Human Communication Technologies Lab[†] Department of Computer Science[‡]

University of British Columbia, Vancouver, BC, Canada

maclean@cs.ubc.ca

ABSTRACT

Visual information overload is a threat to the interpretation of displays presenting large data sets or complex application environments. To combat this problem, researchers have begun to explore how haptic feedback can be used as another means for information transmission. In this paper, we show that people can perceive and accurately process haptically rendered ordinal data while under cognitive workload. We evaluate three haptic models for rendering ordinal data with participants who were performing a taxing visual tracking task. The evaluation demonstrates that information rendered by these models is perceptually available even when users are visually busy. This preliminary research has promising implications for haptic augmentation of visual displays for information visualization.

Categories and Subject Descriptors

H.5.2 [User Interfaces]: Haptic I/O

General Terms

Human Factors.

Keywords

Haptics, 1-DOF, tangible user interface, graspable user interface, haptic perception, multimodal displays, information visualization

1. INTRODUCTION

The visual channel is the primary means for information presentation in many software applications. While this modality is ideal in many situations, there are scenarios where the visual channel can become overloaded by the volume of data being presented. For example, industrial control rooms (e.g. in power plants) are complex visual environments where operators must monitor dozens of readings about various instruments simultaneously. In addition to maintaining an awareness of the state of the plant, operators must interpret and often make decisions about the visual information provided by the instruments. To a lesser extent, we all face visually complex environments where our ability to interpret the entirety of the visual display is compromised: for example, when driving automobiles [11], or while interacting with mobile devices [17].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ICMI '05, October 4–6, 2005, Trento, Italy.

Copyright 2005 ACM 1-59593-028-0/05/0010...\$5.00.

Information visualization attempts to overcome this problem through the use of dynamic presentation or interaction techniques (e.g. focus+context, fisheye distortions and zooming [12]). A complementary approach is to build multimodal displays, where some information from the visual display is offloaded onto a display for another modality, such as touch (c.f. [24] where the approach is to use only the haptic display).

Our focus is on the design of haptic renderings for divided attention contexts where a user is interacting with large datasets. In particular, we aim to haptically convey information rapidly and reliably to users occupied by other visual tasks. These haptic renderings for data should have the following properties:

1. The rendering model should present order information (rendering A is “lower” or “before” rendering B).
2. Renderings should be individually identifiable (rendering A can be identified as “A”).
3. Renderings should be interpretable in visually loaded environments (e.g. control rooms or aircraft cockpits).
4. Renderings should be interpretable by the non-dominant hand.

The first two properties ensure that such rendering models can be used in a wide variety of contexts (namely, those where ordinal data is presented). The third property accommodates the observation that haptic feedback is most likely to be of significant value when other sensory modalities are less available. Finally, the fourth property permits us to conservatively apply our results to either bimanual interaction or to situations where the dominant hand is involved in a different task entirely (e.g. interacting with the graphic user interface). We are interested in haptic renderings for hands since they have the greatest number of nerve endings [10].

Based on these four properties we present a set of haptic rendering models for a 1 degree of freedom (DOF) force-feedback rotary device. At this early stage in our research, the 1-DOF device meets both practical and theoretical needs: the low-cost device was easily accessible; more importantly, the single DOF allowed us to focus on prototyping simple haptic renderings. Because we are interested in these haptic renderings for divided attention contexts, we designed an experiment to evaluate the renderings where participants simultaneously performed a visual tracking task and a haptic identification activity. The results demonstrate that haptic information can be perceived and identified while participants are visually occupied.

Our work makes three primary contributions. First, we rearticulate existing work on haptic feedback within a framework of data types [3], which motivates a rich haptic design space for ordered (ordinal) data. Second, we present a preliminary set of haptic renderings that exploit the ordered property of ordinal data,

and demonstrate their utility in a divided-attention context. Finally, we present a methodology for evaluating haptic renderings that communicate single data points under workload.

The remainder of this paper is organized into four sections. We begin by outlining related work about haptic feedback in divided attention contexts. The review suggests a design space for ordinal data, which has properties that can be exploited by haptic renderings. We then describe our force-feedback device and the development of our haptic rendering models. We follow with a preliminary evaluation of the rendering models, and discuss the implications this work has in the larger context of multi-modal information visualization applications. Finally, we conclude by discussing the implications of our work.

2. RELATED WORK: HAPTIC FEEDBACK

Haptic feedback, which uses the sense of touch as a means for one-way information transmission [16], has been demonstrated to be an effective means to augment other modalities for data transmission (e.g. [21][22]). In our research, we are particularly interested in how and whether this information can be perceived in heavy workload conditions. Within this context, researchers have already considered a wide range of data types for haptic transmission [5]. This body of work covers binary/nominal data, ordinal data, as well as ratio/continuous data (Table 1). We briefly review exemplars of research in each area here.

2.1 Binary/Nominal Data

A binary data point has only one of two values: “on” or “off.” Nominal data is an extension of binary data, where there exist more than two discrete, unordered possible values. A considerable volume of research has examined this type of information transmission, either as a notification system (as in mobile phones), or as a means to cue some other task [14].

For example, tactons are an approach to designing a haptic language for tactile devices [2]. Tactons are sequences of vibratory stimulus combined to convey structured, abstract messages. These vibrations vary on a number of parameters, such as frequency and intensity. For instance, one kind of vibration may represent “File” and another for “Download Complete.” Played one after another, the message would be “File Download Complete.” These messages could be sent to notify the user of system events while the user goes about other graphic user interface (GUI) tasks [2].

This general approach generates haptic renderings with mappings that are iconic in nature. While systematic and highly controllable, the mappings tend to be arbitrary, or metaphoric at best (e.g. [2][15]), therefore restricting the size of the set of haptic renderings based on memory, and restricting the generality of the haptic renderings.

2.2 Ordinal Data

Ordinal data, like nominal data, has discrete values; however, ordinal data values also have an implicit ordering (e.g. musical notes). A growing body of research explores conveying this kind of data, often in car driving contexts, where the driver is considered to be busy with the “driving” task [11]. This research has focused on using haptic-based warning systems to help drivers in simulated driving environments.

Table 1. Different types of data have different properties.

Data Type	Properties	Example
Binary	Two possible values, unrelated	State of a light switch (Boolean)
Nominal	Several discrete values, unrelated	Types of animals (Enumeration)
Ordinal	Discrete values, ordered	Musical scale (Integer)
Ratio	Continuous values	Temperature (Float)

In [11], participants in a driving simulation were given single-level or graded (three-level) collision warnings with either audio or haptic feedback (via vibrations in the chair seat). Participants performed an auditory email sorting task while driving in the simulator. Periodically, a lead vehicle would suddenly brake, potentially causing an accident. The authors found that the graded warnings decreased participant’s reaction times, and were generally preferred over single-level counterparts—especially when the warnings were unreliable (as might be expected in real-world contexts). The haptic warnings were also preferred over the audio warnings.

Beyond driving contexts, presentation of ordinal data has also been of interest in collaborative scenarios as a means of communication. In [3], the authors used a set of vibrotactile output parameters called “hapticons” to facilitate floor control in a distributed groupware application. The application presented a visual workspace, allowing only one user to manipulate the workspace at a time. Collaborators could communicate and perceive the desire for floor control via a vibrotactile interface, using a set of seven different haptic stimuli, divided into a set of three families, each of which were graded (ordinal). Participants learned the set and then were asked to identify hapticons while being occupied by a visual or visual+audio task. Regardless of the type of distraction, users were able to identify the icons through the vibrotactile device with over 95% accuracy.

These results show clear opportunities for haptic presentation of ordinal data, demonstrating that haptics can be effectively used in scenarios where the user’s workload is high. We build on this prior work by generating more explicit ordinal renderings, and by increasing the “graded” nature of the data.

2.3 Ratio Data

Ratio data is continuous, where differences and ratios are interpretable (e.g. distance and weight). Many forms of data in the real world map nicely onto this “continuous” property of ratio data. For example [9] uses a force feedback pedal in a driving simulation to convey the distance to a lead car. This pedal provides ongoing feedback to the driver, thereby giving an implicit warning about when a collision might occur. In contrast to [11], this setup does not require explicit recognition or interpretation of the haptic quantities—which was suitable for the application. Haptic ratio feedback has also been explored as a mechanism to allow blind users to interact with graphs [25]. Using a PHANToM device a physical rendering of the graph was created; although individual points in the data are discrete, the line between points represents a continuous data set and in a typical use case a graph user makes relative judgments regarding




Figure 1. The Twiddler is a force-feedback rotary device.

one value with respect to another. Varying friction coefficients were used to allow users to distinguish between multiple lines on the same graph.

In summary, while several point solutions for building haptic displays under divided attention contexts exist, we do not yet have a systematic understanding of how to build haptic renderings for these environments. For instance, each of the works described above use different haptic displays, varying in form and function. Further, we still do not know the answers to many fundamental questions: how many *different* haptic renderings can people remember? Can memory for these renderings persist longer than a single test session in an experiment? Are the renderings intuitive or general enough to be applied to different contexts?

3. DESIGN OF ORDINAL RENDERINGS

Our interest is in designing intuitive haptic renderings for ordinal data. Our four requirements have driven the design of our haptic models but still leave several open questions. For instance, what is the nature of these renderings? By what process should we build them?

We chose an exploratory approach to design, focusing on ordinal data representations capable of displaying five (5) perceptually distinct values (i.e. very low, low, medium, high, very high)—a large enough range to demonstrate a reasonable level of generality while limiting the complexity of the haptic renderings. Note that we are not seeking to build an optimal rendering yet; we are simply aiming to better understand the problem space.

We begin by describing the capabilities of the one-DOF device we were using to build the renderings. Then, we will describe each rendering in detail.

3.1 Twiddler Device

The Twiddler [20] is a single degree of freedom rotary haptic device (Figure 1). Our Twiddler was equipped with a 20 watt graphite brush DC motor (Maxon RE25 118752), and an HP 4000 line quadrature encoder. In our configuration, the motor was equipped with a plastic knob with a diameter of 65mm. Computer-based control of the renderings took place at a rate of 1000 Hz through a parallel interface. The force bandwidth of the motor was 200Hz.




Figure 2. The detent-wall model represents five levels of information based on the spatial location of the single detent relative to the walls marking the ends of the range. By rotating the knob counter clockwise, the user (triangle) can assess the distance to the “low” wall (and vice versa for the “high” wall).

3.2 Haptic Rendering Space for Ordinal Data

We evaluated three different models from two general classes of haptic renderings: “position based” renderings which require user-driven exploration through the rendered space to discern the information, and “time based” renderings which actively transmit haptic information to the user (i.e. hapticons). We were interested in exploring both rendering types, because while users can more carefully explore position based renderings, the time based renderings can potentially convey information more rapidly and in a more controlled manner.

When designing the renderings, we had three requirements.

- The user should be able to identify the data point with minimal movement; a full turn of the knob (360°) was considered the upper limit.
- Where possible, the rendering should exploit the ordered property of the ordinal data type.
- Where possible, we should use natural perceptual or motor mappings.

Rendering ordinal data offers one distinct benefit over nominal data: namely, ordinal data is ordered. Properly designed, rendered data points should also seem sequentially connected rather than an arbitrary collection of signals. As a result they should be easier to learn since the understanding of one signal applies to other signals.

With the 1-DOF Twiddler device, we explored a number of variables for rendering the ordinal information. We considered parameters such as rotational displacement (or location), torque (or applied force), amplitude and frequency of vibration, as well as resistance and damping. These explorations led us to four models which seemed promising for ordinal information display.

We do not believe that this set of renderings is in any respect exhaustive. However, the renderings we describe include representatives on the several axes of the potential design space, including mappings based on location, differential force, and vibratory frequency.

3.3 Position Based Renderings

We developed three different position based renderings for haptically displaying ordinal data. These renderings met the properties outlined above, and two were selected and refined prior to final evaluation based on informal non-workload pilot studies.




Figure 3. The torque-differential model renders different torques for each rotational direction (top). The torque always pushes back towards top center; the difference between the clockwise and counterclockwise torques conveys the level. The narrow wedge of zero torque in the center facilitates stable rendering.

3.3.1 Detent-wall model

The detent-wall model was based on a spatial representation of the data value (Figure 2). The entire region was mapped to a 270° turn on the knob¹. The extents were represented by walls rendered using an exponential function that exerted the maximum sustained force available through our amplifier and motor. A single detent, representing the data value, was rendered at one of 45° (Very Low), 90° (Low), 135° (Medium), 180° (High), and 225° (High).

3.3.2 Torque-differential model

The torque-differential model was based on the literature which suggests that humans are better at discerning *differences* between haptic stimuli than absolute values [18]. The torque-differential model presents two *different* directional torques depending on the direction in which the knob is turned: when the knob is turned clockwise, the torque is applied counter-clockwise, and vice-versa (Figure 3). The difference between these torques represents each of the levels on the scale (Figure 3).

3.3.3 Ramp model

The ramp model was also based on a physical model and the literature which suggests that users are good at discerning differences between stimuli. This rendering modeled turning a knob up two parabolic ramps. In this case, instead of discerning the torque being exerted against the knob in either direction, the user would detect differences in the torque's *rate of change*. Since the stiffness would be greater on one side, we believed this model had the potential to perform better than the torque-differential model [22]. Pilot testing yielded very poor results both in time and accuracy with the ramp model; consequently, we discarded this rendering from our final experiment.

3.4 Time Based Rendering

We used simple sinusoidal hapticons to render level information. This rendering mapped the speed of the vibration to the level. Using [15] as a guide, we rendered sinusoidal force waves of 4-20Hz, and an amplitude of 0.5 (Table 2). In this rendering model, the user passively experienced force sinusoids independent of knob position. Although we used a force feedback device to enable comparisons with our other renderings, in practice, comparable renderings could be generated with a vibratory display.

¹ 270° was chosen for the detent-wall model through empirical pilot testing to reduce the degree separation between the extents while allowing sufficient resolution.

Table 2. Hapticon sinusoidal frequencies for each level.

Level	Frequency
Very Low	4 Hz
Low	8 Hz
Medium	12 Hz
High	16 Hz
Very High	20 Hz

4. EXPERIMENT

Our experiment was designed to evaluate the effectiveness and efficiency of the three different ordinal rendering models in the presence of a visually demanding task. We had two specific goals: first, to understand whether the ordinal haptic information could be perceived and interpreted under a visually taxing task; secondly, to determine which renderings would yield the highest accuracy and fastest response.

4.1 Hypotheses

We had three hypotheses for this study. Because the time based rendering transmitted information without user intervention, we expected this rendering to have the fastest response time.

Hypothesis 1: The time based rendering elicits faster response times than other models.

However, due to the challenge of perceptually resolving vibration frequencies absolutely [16], we also expected the time-based rendering to have the lowest accuracy.

Hypothesis 2: The time based rendering elicits poorer accuracy than the other models.

Thirdly, because the torque-differential model requires smaller motions on the part of the user to assess a value, we expected this model to outperform the detent-wall model

Hypothesis 3: The torque-differential model will have faster performance than the detent-wall model.

4.2 Participants

A total of fifteen (15) paid computer science or engineering graduate students participated in a one-hour session. Participants ranged in age from 20-41. Eleven participants were male, and four were female. All participants but one were right-hand dominant. Nine of the students had prior experience with computer-driven haptic devices (participants in prior studies at the university), but none were familiar with our work. None of the participants who had taken part in an earlier pilot study to guide the design of our models participated in this experimental evaluation.

4.3 Apparatus

The experiment was conducted using custom-built Windows XP applications running on a Pentium 4 – 2GHz and the Twiddler hardware described above. The GUI was displayed on a 17" LCD monitor at 1280×1024 resolution. Participants interacted with the applications through a standard mouse placed near their dominant hands, and the Twiddler device placed near their non-dominant hands (Figure 4).

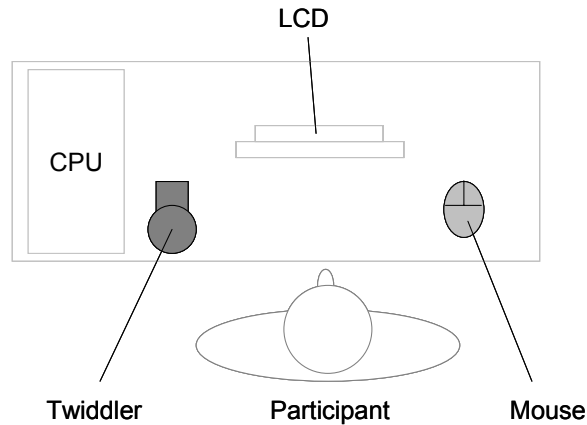


Figure 4. The experimental setup (cables not shown). The mouse was placed on the dominant-hand side, and the Twiddler on the non-dominant side.

4.4 Tasks

Participants completed trials on a divided attention task where both tasks are performed simultaneously (Figure 5). The first task involved multi-target visual tracking and the second task was a haptic assessment task.

The visual tracking task required the participant to track three out of six moving objects in a separate window, a task that has been shown to be cognitively taxing [19]. When prompted at 60s intervals, participants were required to identify the three tracked objects by clicking on them with the mouse. Following selection and commitment, the three correct objects would blink briefly to indicate which objects were to be tracked for the next trial.

In the haptic assessment task, haptic renderings were displayed on the device. Participants were welcome to interact with the device in any way they chose. Most interacted through grasping, although some participants also tried using a single finger to rotate the device. Participants were then asked to identify the data value that was being presented (very low, low, medium, high, or very high) by pressing a corresponding button on the GUI interface (Figure 5). Once the trial was complete, the next haptic rendering would be displayed following a brief pause (10s) and an audio beep.

Note that both tasks were performed simultaneously so the timing of the haptic presentation and the visual prompts were asynchronous. While we were primarily interested in speed and accuracy on the haptic assessment task, we also needed to assess performance on the visual task as a measure of the cognitive load actually experienced by participants.

4.5 Experimental Design

In a within-subjects design, participants completed a total of twenty (20) trials with each of the three haptic renderings, blocked and counterbalanced by rendering. These trials yielded four replicates of each level for every rendering, displayed in random ordering within a rendering block. The experiment was designed so that for each participant, we could collect response time for each trial and aggregate accuracy data for each rendering model. For the visual task we collected both response time and accuracy for each trial. The experiment was designed so that roughly five haptic trials per model would be interrupted by the

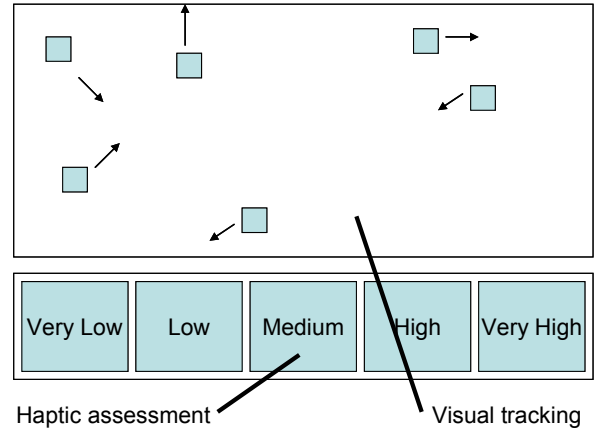


Figure 5. A mock-up of the software used in the evaluation, showing the haptic assessment task (bottom) and the visual tracking task (top) running concurrently.

visual task; speed data for these trials was discarded due to the unfair “double penalty” inflicted on slower haptic models of having more visual task interruptions.

4.6 Procedure

Participants completed a pre-test questionnaire to collect demographic information such as experience with computer-based haptic and graphic user interfaces.

Participants then completed trials of the haptic assessment and visual tracking task. Haptic assessment trials displayed one of five levels in a random order, counterbalanced by rendering. Prior to beginning each rendering block, participants completed a training period with that rendering which, and had to complete a baseline recognition test for each level with at least 90% accuracy.

Once all trials were complete, participants completed a post-test questionnaire to assess subjective preference information. The participants were subsequently debriefed regarding the research objectives of the experiment.

5. RESULTS

The results for the first participant were discarded because during the experiment, it became clear the participant had not understood the experimental procedure. To prevent further misunderstandings, we revised the training procedure for the remaining participants. In total then, 1260 trials (14 participants \times 30 trials \times 3 rendering types) of haptic assessment tasks were completed. Of these, we discarded 143 data points since the haptic trials were interrupted by the visual task, thereby inflating the response time. In total, 208 trials of the visual task were completed (note that not all visual task trials interrupted haptic trials).

Haptic trial times were aggregated by model on a per-user basis (Table 3). The results show that participants completed trials more quickly using the hapticon (\bar{X} = 3.24s, σ^2 = 0.66) and torque-differential (\bar{X} = 4.40s, σ^2 = 1.32) renderings compared to the detent-wall model (\bar{X} = 6.02s, σ^2 = 1.61).

We set alpha at a 0.05 level. A one-way repeated-measures ANOVA confirmed that one haptic rendering was faster than the others ($F_{2,39}=17.6$, $p<0.05$). A subsequent Tukey post-hoc analysis revealed that the hapticon model was significantly faster than both the torque-differential ($t_{13}=4.54$, $p<0.016$) and the detent-wall model ($t_{13}=7.60$, $p<0.016$). Further, it revealed that the torque-differential model was faster than the detent-wall model ($t_{13}=6.01$, $p<0.16$). These results support hypotheses 1 and 3, namely that the hapticon model was faster than the other two renderings and that the torque-differential model was faster than the detent-wall model.

Participants were quite accurate in determining the level being rendered regardless of the model (Table 3). For the detent-wall model participants averaged 83.9% correct ($\sigma^2=17.9$). Participants performed most accurately with the torque-differential model at 92.9% ($\sigma^2=8.1\%$) and least accurately with the hapticon model at 74.6% ($\sigma^2=15.6\%$). Again setting alpha at a 0.05, a one-way repeated measures ANOVA supports our hypothesis that at least one rendering was less accurate than the others ($F_{2,39}=5.67$, $p<0.05$). A Bonferroni post-hoc analysis revealed that the hapticon model was less accurate than the torque-differential ($t_{13}=4.39$, $p<0.025$), but no less accurate than the detent-wall model ($t_{13}=1.82$, $p=0.045$). This result supports hypothesis 2 which states that the hapticon model is less accurate than the other models.

Participants also performed extremely well with the visual tracking task. Out of 208 trials participants accurately tracked all three objects in 166 trials, and two of three objects in 196 trials with no meaningful effect for model-type. This result indicates that participants were actively engaged with the visual task. Results of the post-study questionnaire confirm that subjectively, participants felt that the visual task occupied most of their attention (13/14 participants indicated that it took up “Most” of their attention while 1 participant indicated that it took up “All” of his attention).

With novel interfaces it is valuable to consider subjective perception of efficiency and accuracy. When participants were asked to rank-order the haptic models based on how confident they were of their responses, the torque-differential model received 10 top votes. In contrast the hapticon approach received 11 votes for producing low-confidence ratings. This perception parallels the actual performance of the renderings (as above). With alpha set to 0.05 a chi-square analysis supports the hypothesis that participants favored the torque-differential model. Interestingly while 8/13 participants indicated that they believed that the visual task impaired their judgments of the haptic information (marking 3 or higher on a 4 point scale), participants still performed quite admirably (83.8% across all haptic renderings).

When asked to rank-order the haptic models in terms of how quickly they were able to retrieve information, the torque-differential model received 9 votes for being quickest while the hapticon model only received 3 votes. Interestingly, the hapticon approach received 9 of the votes for being the slowest, in contrast to the quantitative data showing that the hapticon rendering was fastest. However, this perception of “quickness” may have been related to how challenging participants found each rendering model was to use—previous work has demonstrated that when

Table 3. Performance and accuracy of each model (n=14).

	Detent-wall	Torque-differential	Hapticon
Time per trial	6.02s	4.40s	3.24s
Std Dev time per trial	1.61s	1.32s	0.66s
Accuracy	83.9%	92.9%	74.6%
Std Dev accuracy	17.9%	8.1%	15.6%

participants have difficulty with interfaces, the perceived duration of use is longer [7]

Ultimately participants were able to perceive and accurately interpret the ordinal information presented on the haptic interface regardless of the rendering type.

6. DISCUSSION

The results of our experiment are promising for three reasons: first, participants were able to perceive and interpret haptic information while their visual attention was occupied; secondly, participants were able to interpret this information rendered on a 1-DOF device held in their non-dominant hands, and finally, these renderings achieved levels of accuracy acceptable in many applied contexts. We envision at least two ways in which renderings like these could be used in information visualization: focus+context displays and divided attention contexts.

Many focus+context techniques often render large data sets with significant visual distortion [12]. A multi-modal approach, rendering focus information on a haptic display and context information on the visual display (or vice versa) could sidestep the problems of distortion-based renderings [3] or the physical real-estate problems of large-display approaches [2]. Promising research in a similar vein explores the use of peripheral vision to encode “extra” data to augment the focal information [1]. Peripheral vision has significantly lower resolution than the region under the fovea; the results indicate that peripheral vision, a low bandwidth channel, is capable of conveying useful information in parallel with reading text on a display. This supports our proposition that applications of refined versions of our rendering models could be used, for example, to provide coarse grained level information in an industrial plant (e.g. efficiency data).

The focus+context approach would only be applicable in certain situations. For example, in an industrial plant, time-series vibration data from individual sensors which requires a sizable visualization could be rendered on the haptic display while an operator interacts with an interactive overview GUI “map” of bearing vibration levels. Another area that could be explored involves using a haptically enabled mobile device, such as a PDA, where the environment itself provides context, while visual and haptic details are provided on the device. This might lend itself particularly to applications where touching the physical environment can yield useful information but may be undesirable or dangerous such as in waste processing or industrial contexts.

In divided attention contexts, haptic information can be used to convey divergent environmental properties which may be important to the user. Examples of this include operating a motor

vehicle or aircraft (e.g. [8][9][11][21]). Haptic monitoring devices are another example of a divided attention application of our results. Such devices would be able to operate at the periphery of the user's attention without distracting from the primary task. Additionally, applications could use haptic displays as secondary displays, either by augmenting information presented to the visual modality (as in [9]), or by presenting supplementary information that is related to but not present on the visual display [13]. In this latter scenario, the user needs to be able to perceive information from the haptic display and to make decisions about that information while potentially interacting with a visual display. Although the fundamentals of this divided-attention task have been shown in the past for explicit low bandwidth level information (e.g. [3][14][22]), we have extended these findings apply for ordinal data on 1-DOF devices in the non-dominant hand.

Our current program of research is geared toward developing these multimodal displays (i.e. using both visual and haptic displays). In applications with such displays, we expect that when related information is presented on the different displays, interpreting information from the haptic display should be easier than in our experimental context. Anticipating that the divided attention task would be a stumbling block to applied uses of these multimodal displays, we designed our experiment and renderings to understand the divided attention problem. The positive results of our experiment with unrelated haptic/visual stimuli suggest that even in a worst case scenario, the multimodal approach is viable.

7. CONCLUSION AND FUTURE WORK

In summary, we have developed a set of rendering models (detent-wall, torque-differential, hapticon) that fulfill our four design requirements: they are all capable of rendering ordered information while being individually identifiable; the renderings are interpretable given a taxing visual tracking task and finally, the renderings were all usable with the non-dominant hand. These results demonstrate that the haptic channel is a viable means of explicit information transfer while the user's visual channel is heavily loaded.

We were also interested in the very specific context of ordinal data where individual items need to be identifiable (although there exists a wide variety of domains where this requirement can be relaxed). Overall, the high level of accuracy with all three models suggests that iterative improvements could make each one a component of a larger haptic vocabulary for 1-DOF devices; these improvements will be considered in future work. For example, the hapticons could be designed to be more perceptually distinctive (e.g. by using the Weber-Fechner Law to develop a perceptually linear scale). A more carefully chosen set of hapticons could potentially produce stronger results for this model (e.g. [3]).

The rendering models we presented take advantage of the ordered property of ordinal data. In the detent-wall model, this ordering was mapped spatially. For the torque-differential model, we mapped this ordering property to perceptual differences by applying different forces to the user depending on which way the knob was turned. Finally, the hapticon model applied the mapping property literally to different frequency vibrations. Although we did not test this hypothesis, comments from the participants suggest that they used these mappings, rather than muscle memory, to identify the haptic levels, for example:

P: "With [the detent-wall model], all I had to do was figure out how far I was from the wall. Sometimes, I'd guess the closer wall, so I'd get there faster."

Comments of this type are promising for three reasons. First, if we can design models that exploit psychophysical "sweet spots," then such renderings should be easier to learn, and easier to recall. Second, because the mappings we use have a natural progression, it should be possible to introduce additional intermediary "levels" to the models. Finally, because the mappings are not iconic, but instead map to abstract "levels," they can be used in a wide variety of contexts.

Using a divided-attention task, we have demonstrated that humans are able to perceive and interpret interval data presented on a haptic, 1-DOF display. We hope to expand this body of knowledge by exploring other forms of data (e.g. continuous, time-series) to map out the "haptic rendering space" for 1-DOF devices. With this knowledge, we would be able to effectively build multimodal information visualization displays for a variety of data types and tasks [12].

8. ACKNOWLEDGMENTS

We thank Warren Cheung, Mario Enriquez, Joseph Luk, Adam Bodnar, and Sid Fels for insightful discussions on this work, and Colin Swindells for his technical expertise.

9. REFERENCES

- [1] Bahna, E., Jacob, J. K. Augmented reading: presenting additional information without penalty. In *Extended abstracts of the SIGCHI conference on Human factors in computing systems (CHI '05)* (Portland, USA, April 2-7, 2005). ACM Press, New York, 2005, 1909-1912.
- [2] Baudisch, P., Good, N., Bellotti, V., and Schraedley, P. Keeping things in context: a comparative evaluation of focus plus context screens, overviews, and zooming. In *Proceedings of the SIGCHI conference on Human factors in computing systems (CHI '02)* (Minneapolis, USA, April 2002). ACM Press, New York, NY, 2002, 259-266.
- [3] Brown, L., Brewster, S., and Purchase, H. A first investigation into the effectiveness of tactons. In *Proceedings of the first joint eurohaptics conference and symposium on haptic interfaces for virtual environment and teleoperator systems (WHC 2005)* (Pisa, Italy, March 18-20, 2005), 2005.
- [4] Carpendale, M.S.T. Elastic presentation space. *Information Design Journal*, 10(1), 2001, 58-69.
- [5] Card, S. and Mackinlay, J. The structure of the information visualization design space. In *Proceedings of the 1997 IEEE Symposium on Information Visualization (InfoVis '97)* (Washington, USA), 1997.
- [6] Chan, A., MacLean, K., and McGrenere, J. Learning and identifying haptic icons under workload. In *Proceedings of the first joint eurohaptics conference and symposium on haptic interfaces for virtual environment and teleoperator systems (WHC 2005)* (Pisa, Italy, March 18-20, 2005), 2005.
- [7] Czerwinski, M., Horvitz, E., and Cutrell, E. Subjective duration assessment: an implicit probe for software usability.

- In *Proceedings of IHM-HCI 2001 Conference* (Lille, France, 2001), 167-170.
- [8] Enriquez, M., Afonin, O., Yager, B., and Maclean, K. A pneumatic tactile alerting system for the driving environment. In *Proceedings of the 2001 workshop on Perceptive user interfaces*. (Orlando, USA, November 15-16, 2001). ACM Press, New York, NY, 2001, 1-7.
- [9] Enriquez, M., and MacLean, K. Impact of haptic warning signal reliability in a time-and-safety-critical task. In *Proceedings of the 12th annual symposium on Haptic interfaces for virtual environments and teleoperator systems (IEEE-VR2004)* (Chicago, USA, March 27-31, 2004), 2004.
- [10] Kandel, Eric R., Schwartz, James H., and Jessell, Thomas M. *Principles of Neural Science*, 4th Ed. McGraw-Hill Companies, Inc., 2000.
- [11] Lee, J., Hoffman, J., and Hayes, E. Collision warning design to mitigate driver distraction. In *Proceedings of the SIGCHI conference on Human factors in computing systems (CHI '04)* (Vienna, Austria, April 24-29, 2004). ACM Press, New York, NY, 2004, 65-72.
- [12] Leung, Y., and Apperley, M. A review and taxonomy of distortion-oriented presentation techniques. *ACM Transactions on Computer-Human Interaction*, 1(2), 1994, 126-160.
- [13] Lindeman, R., Sibert, J., Mendez-Mendez, E., Patil, S., and Phifer, D. Effectiveness of directional vibrotactile cuing on a building-clearing task. In *Proceedings of the SIGCHI conference on Human factors in computing systems (CHI '05)* (Portland, USA, April 2-7, 2005). ACM Press, New York, 2005, 233-444.
- [14] Lindeman, R., Yanagida, Y., Sibert, J., and Lavine, R. Effective vibrotactile cuing in a visual search task. In *Proceedings of the ninth IFIP TC13 international conference on Human-computer interaction (INTERACT 2003)* (Zuerich, Switzerland, September 1-5, 2003). ACM Press, New York, NY, 2003, 89-96.
- [15] Maclean, K., and Enriquez, M. Perceptual design of haptic icons. In *Proceedings of EuroHaptics 2003* (Dublin, Ireland, July 6-9, 2003), 2003.
- [16] MacLean, K. Designing with haptic feedback. In *Proceedings of IEEE robotics and automation (ICRA 2000)* (San Francisco, USA, April 22-28, 2000), 2000.
- [17] Oulasvirta, A., Tamminen, S., Roto, V., and Kuorelahti, J. Interaction in 4-second bursts: the fragmented nature of attentional resources in mobile hci. In *Proceedings of the SIGCHI conference on Human factors in computing systems (CHI '05)* (Portland, USA, April 2-7, 2005). ACM Press, New York, NY, 2005, 23-45.
- [18] Pang, X., Tan, H., and Durlach, N. Manual discrimination of force using active finger motion. In *Perception & Psychophysics*, 49(6), 1991, 531-540.
- [19] Pylyshyn, Z., and Storm, R. Tracking multiple independent targets: evidence for a parallel tracking mechanism. In *Spatial Vision*, 3, 179-197.
- [20] Shaver, M., and MacLean, K.M. The twiddler: a haptic teaching tool design manual: low cost communication and mechanical design. Report TR-2005-09, Department of Computer Science, University of British Columbia, Vancouver, Canada, 2005.
- [21] Sklar, A., and Sarter, N.. "Good vibrations": the use of tactile feedback in support of mode awareness on advanced technology aircraft. *Human Factors*, 41(4), 1999, 543-552.
- [22] Tan, H., Durlach, N., Beaugard, G., and Srinivasan, M. Manual discrimination of compliance using active pinch grasp: The roles of force and work cues. *Perception & Psychophysics*, 57(4), 1995, 495-510.
- [23] Vitense, H., Jacko, J., and Emery V. Multimodal feedback: establishing a performance baseline for improved access by individuals with visual impairments. In *Proceedings of the fifth international ACM conference on Assistive technologies (ASSETS 2002)* (Edinburgh, Scotland, July 10-12, 2002). ACM Press, New York, NY, 2002, 49-56.
- [24] Wall, S. and Brewster, S. Application based assessment of frictional properties for haptic data visualization. In *Proceedings of EuroHaptics 2004* (Munich, Germany, June 5-7, 2004). 2004.
- [25] Yu, W., Ramloll, R. and Brewster S.A. Haptic graphs for blind computer users. In *Proceedings of the First Workshop on Haptic Human-Computer Interaction*, 2000, pp 102-107.