2019 IEEE World Haptics Conference (WHC)
Tokyo, Japan
9-12 July 2019

Single-Actuator Vibrotactile Numeric Information Delivery in the Face of Distraction

Jeffrey R. Blum and Jeremy R. Cooperstock

***Abstract*— We test ActiVibe, a previously reported method for communicating numeric values between 1 and 10, to determine whether it remains optimal under conditions reflective of more challenging potential real-world use cases. We thus consider vibrotactile communication in conjunction with an audio distractor task, and when conveying not just one, but three numeric values in succession. Results of a user study comparing three different rendering methods indicate that ActiVibe maintains both accuracy and subjective preference advantages vs. two different duration-based methods when conveying a single value, but largely loses these advantages when presenting three sequential values. Under conditions similar to the most difficult ones we test, a more concise duration-only approach may be preferable for some applications, requiring less power consumption and demanding attention for less total time.**

I. Introduction

There is an increasing desire to use haptics to render information beyond simple alerts on mobile devices, such as conveying progress toward a goal without requiring a screen. With users wearing devices in daily life, designers must make haptic information display robust in real-world use.

The work presented here evaluates Cauchard et al.'s ActiVibe [2], which successfully conveys the numbers 1 to 10 through vibrotactile Tactons [3]. ActiVibe uses a sequence of vibration pulses that the user can count. An in-the-wild study showed that participants could accurately perceive infrequent single values throughout the day even when going about their daily routines. However, we hypothesized that ActiVibe's multiple pulses for each value could become difficult to track in cases such as when the user is highly distracted or receiving multiple values sequentially in a single message.

Such concerns may be particularly relevant for applications that provide a frequent, periodic, multi-value background indicator, such as Blum and Cooperstock's SenseProxy [4]. SenseProxy uses a series of vibration pulses whose durations correlate to parameters of a remote partner, such as amount of leg motion and overall velocity. We propose that for such applications, a simpler, non-counting pattern may be more robust in the presence of a simultaneous audio task, and still sufficiently accurate for background status updates.

Our main contribution is in demonstrating that a Tacton design optimized for accuracy in comparatively simple conditions loses its advantages under more difficult conditions, vs. a duration-only approach. Specifically, we evaluate:

1) two of Cauchard et al.'s ActiVibe designs in the presence of a foreground audio task.
2) a novel "Baseline PreVibe" (BPV) rendering that provides a perceptual calibration stimulus before information vibration(s), tested against both ActiVibe options.
3) all three designs conveying three values in succession.

II. Background and Related Work

Haptic feedback, although traditionally used for simple notifications, can convey information directly, with the goal of reducing or eliminating the need to look at a screen when receiving simple messages. This has been successful in the form of Tactons [3], or haptic icons. Especially since it does not interfere with visual or audio channels, haptic feedback is particularly suited to delivering background information [5]. For example, Brewster and King used Tactons to present progress for a long-running process, finding that a multi-part sequence of vibrations performed better at providing background status than a visual progress bar [6]. More generally, Pielot and de Oliveira showed that a perceptually near-threshold ongoing vibration can fade into the background, and yet be noticed when it stops [7].

Distraction from competing tasks can can be particularly disruptive when attempting to deliver information in the background. Oakley and Park tested three different distractor tasks (transcribing poems, data-entry, and walking) while participants received Tactons that varied by their location around the wrist and roughness, finding that distraction significantly lowered the recognition rate [8].

Thus, creating Tactons that are succinct yet clear even when the receiver is engaged in another task, is critical for real-world background haptic information delivery. Saket et al. used two different vibration durations interspersed with two different gap lengths to convey urgency of phone calls, finding that the four levels they tested were differentiable [9]. Brown et al. tested varying rhythm, roughness and the location on the body to represent three different pieces of information in a single Tacton, finding that the number of levels is crucial, e.g., restricting to only two levels of roughness performed much better than three [10]. Tang et al. explored using a rotary dial with different tactile encodings to convey five ordinal levels during a visual distractor task [11], a goal similar to the work presented here. They found that active exploration by twisting the dial, rather than passively receiving a frequency-varied hapticon, performed best. This active approach may be useful in some interfaces, but not

This work is an expanded version of an extended abstract presented as a poster at Ubicomp 2018 [1]. This work was supported by an NSERC Discovery Grant, a McGill Graduate Mobility Award, and a Mitacs Globalink Research Award. McGill REB #83-0814.

Both authors are from the Shared Reality Lab, Dept. of Electrical and Computer Engineering and the Centre for Interdisciplinary Research in Music Media and Technology (CIRMMT), McGill University, Montreal, Quebec, Canada. jeffrey.blum@mail.mcgill.ca, jer@cim.mcgill.ca

978-1-5386-9461-9/19/$31.00 © 2019 IEEE

for background information delivery. Wang et al. found a limit of approximately 4 or 5 levels of vibration, separated by frequency of the stimulus, that could be detected [12].

Numerosity, or the ability to count sequential tactile, visual or audible pulses [13] has been shown to be effective in practical tasks [14], as well as explored in multisensory situations [15]. Numerosity was also used by Pasquero et al., who found that rendering countable pulses via a piezo-electric actuator in a wristwatch worked best at around 3 pulses/second [16]. Slower rates made the stimulus annoyingly long, and speeding up caused more errors. They tested their rendering in an office environment for informing the wearer of their unread email count, and anecdotally report that some participants tried it while conversing with co-workers, but no formal analysis was done concerning using the system while distracted. In their study, the user explicitly instigated the rendering each time they wanted to know their email status, so although subtle, it was not tested for background, ongoing information delivery.

III. ActiVibe and Baseline PreVibe

The ActiVibe final design (AVF) [2] uses easily countable discrete vibrations, augmented with longer actuations representing "5" (Figure 2a). Although ActiVibe overall performed well in the wild, participants only responded to 79% of the stimuli, and accuracy while very active (e.g., running) dropped to 54%. Participants were "Stationary" during 67% of trials, which may correlate with being relatively undistracted. Thus, it is unclear from the ActiVibe study how much physical masking of the signal from activity, vs. differences in mental distraction, impacts performance. Audio distraction is also addressed by Cauchard et al., who cite research indicating that attending to speech makes it particularly difficult to perform numerical tasks [17], as ActiVibe demands. Although they suggest detecting people talking nearby, presumably to time the vibrations to occur during silent periods, we hypothesized that removing the explicit counting requirement may also help to ameliorate the problem. We were therefore motivated to test under controlled, more strenuous conditions, as well as try a novel rendering. We use a verbal audio distractor task (Section IV-B), rather than a physical confound such as exercising or fidgeting, to focus on the effects of mental attention on a main task. An alternative distractor, such as physical movement, a visual tracking task, a separate haptic task, or a non-verbal audio task, may all impact the results found in this study, and are compelling directions for future work.

Cauchard et al. explored six patterns. The worst-performing was a duration-only (DO) method (ActiVibe "B"), with a single vibration lasting 100 ms times the value (e.g., the value "3" is $3 \times 100 = 300$ ms). However, DO was the only option not requiring counting multiple vibrations per value, making it potentially more robust under load or when rendering multiple values (Figure 2b). The mean DO error magnitude ranged between approximately 0.5 and 2, with larger target values exhibiting worse accuracy. This leads to two observations. First, although accuracy is worse with DO than AVF, it performs adequately if only an approximate value is needed. Second, improved performance at higher values will have the largest impact on overall accuracy, since DO error is lower when conveying values at the bottom of the range, and consistently larger for higher values.

These observations, coupled with reported user desire for a "heads-up" introductory PreVibe during the ActiVibe study [2], lead us to propose using a PreVibe to not only focus attention, but also to use its duration to represent a baseline value. This allows the user to compare the PreVibe duration to a following DO information pulse to calibrate its meaning. We refer to this new design as *Baseline PreVibe* (BPV, Figure 2c). With BPV, the pattern begins with a 500 ms vibration that reminds the user what a value of 5, or 50% of the range, feels like. For example, a following vibration longer than the PreVibe is a value greater than 5. A vibration twice the PreVibe duration is 10.

We hypothesized that BPV would have greater accuracy vs. the pure DO design, yet also require less mental load, and therefore be more resistant to distraction, than the counting-based AVF approach. BPV may be particularly beneficial for higher values, since the PreVibe anchors the middle of the range, providing a baseline at a value of 5, from which higher values may be more readily ascertained.

IV. Experiment Design

The experiment had two sequential phases. Phase 1 tested a single value per trial, and Phase 2 tested three values per trial. Before describing the procedure, we first describe the haptic perception and audio distractor tasks.

A. Haptic value perception task

Participants sat at a table and rested their non-dominiant arm on a towel while wearing a Pebble model 301 smartwatch. They used a tablet interface (Figure 1) to enter the value(s) they perceived. Although not done in the ActiVibe lab studies, participants also wore headphones playing pink noise to mask the audible Pebble vibration motor. Participants were instructed to enter their best guess even when unsure, but to use the "?" (default response) at the bottom of each column if they had no idea of the value presented. Values could be changed until the Submit button was tapped. Trials continued automatically with 8 s between the end of the vibrations from one trial and the beginning of the next.

The columns and Submit button disappeared when the Submit button was pressed, and reappeared at the very end of each trial, just before the vibrations stopped. This precluded participants from entering the values one at a time, immediately after feeling each value. In practical use, we expect that users would typically need to process information from all of the values presented as a whole, so we believe this increased ecological validity.

B. Simultaneous audio task (Blues)

To simulate real-world distracted use, when receiving haptic information is a background task subordinate to a main task such as driving a car or working, participants were given

Abb.	Name	Brief Description	PreVibe (ms) [gap duration]	Value (ms) [gap duration]	Average total pattern time (ms)	Average time vibe motor active (ms)
AVF	ActiVibe-Final design	Roman-numeral style	700 [1200]	value pattern [200]	2960	1420
BPV	Baseline PreVibe	"5" calibration + DO	500 [1200]	value*100	2250	1050
DO	Duration Only	ActiVibe design B	none	value*100	550	550

TABLE I: Rendering conditions, with abbreviations as used throughout paper. Note that only DO does not have a PreVibe. Durations in [] represent gaps, or pauses between vibration pulses. The vibration durations for AVF are identical to the ActiVibe longitudinal study: 150 ms short pulse, 600 ms long (value 5) pulse, 200 ms between pulses. The last two columns provide the mean pattern duration for a single value (averaged across the values one through ten), including both vibration time and gaps, and the mean duration the vibration motor is actually running, correlated with power consumption.

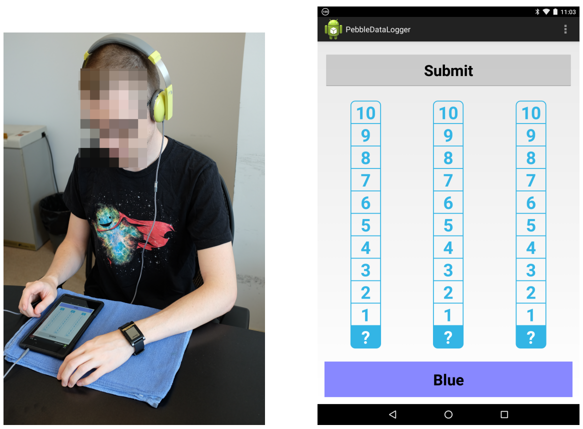

Fig. 1: Experiment setup and tablet UI for entering responses.

an audio task, similar to that used by Chan et al. [18]. The Android tablet's text-to-speech engine (ASUS Nexus 7 2013) spoke color names over the pink noise. Participants were instructed that their most important task was to tap the Blue button (Figure 1) as quickly as possible each time "blue" was spoken. Blue was spoken four times in a randomized list along with 16 other color names, one color per second, after which the list was re-randomized. Thus, four blue stimuli occurred every 20 s, or 20% of the time.

C. Experiment procedure

Twelve participants (plus one removed as detailed below) were recruited from the University community (7 male, 5 female; ages 21–37, median=24) and compensated CAD$15. The experiment took approximately 80 minutes. After consent and a pre-questionnaire, the participant listened to the experimenter read from a script describing the Blue audio task and explaining they would simultaneously be receiving smartwatch vibrations that represented different numbers.

The Pebble watch was strapped comfortably tightly to the participant's non-dominant wrist, leaving their dominant side unencumbered for manipulating the tablet UI. Participants briefly practised tapping the Blue button when hearing blue spoken in the headphones, while entering values based on the number of fingers held up by the experimenter near the tablet screen. Once it was clear that the participant was using the interface correctly, Phase 1 began.

1) Experiment Phase 1 (One Value): Phase 1 compared all three renderings (BPV, DO, AVF) when conveying a single value. Participants first received a scripted verbal description of the current rendering pattern while referencing a printed graphical representation (Figure 2). Timings for the rendering conditions can be found in Table I. Next, the participant was exposed to all ten values, in order, to familiarize them with the rendering. The values about to be rendered were displayed as numbers on the watch for training purposes.

Each of the six possible orderings of the three rendering conditions was presented to two of the 12 participants. Within each rendering condition, all ten values were presented in random order, three times, resulting in 30 trials per rendering condition. After completing all three rendering conditions, participants filled out a questionnaire, took a short break, and proceeded to Phase 2.

2) Experiment Phase 2 (Three Values): Phase 2 compared the same three rendering conditions when conveying not just one value, but rather three values sequentially in each trial, referred to as BPV3, AVF3, and DO3. As in Phase 1, the six possible orderings of the three rendering conditions were each presented to two of the 12 participants. The order of the six possible rendering conditions in Phase 2 was independent of that used in Phase 1. Participants received ten training trials, with each possible value presented once in each of the three positions, in random order (e.g., 5-1-7, 8-5-2, etc.).

PreVibe and Tacton parameters for individual values were identical to Phase 1. Only a single PreVibe was delivered before each set of three values for AVF and BPV. A PreVibe before each individual value would likely be very confusing, and particularly penalize AVF by making it more difficult to keep one's place in the overall pattern. An 800 ms gap between values was determined via self-experimentation and pilot testing, which is 4 times the gap between vibrations within a single AVF value, and 400 ms shorter than the 1200 ms gap between the PreVibe and values in AVF and BPV. Before each rendering condition, the experimenter verbally repeated key parts of the rendering description. After all three rendering conditions, the same questionnaire from Phase 1 was given.

V. Results

A. Data cleaning

Participant p08 was rejected for not following instructions, and was replaced by p13. Participants p03 and p11 restarted one condition each due to a technical issue that also required removal of spurious audio stimulus entries in p11's log file.

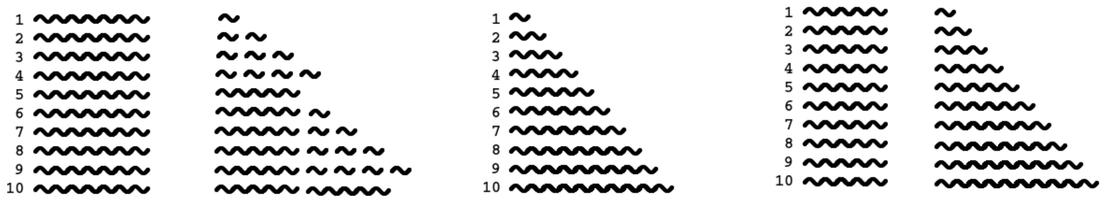

(a) AVF (ActiVibe Final) (b) DO (Duration Only) (c) BPV (Baseline PreVibe)

Fig. 2: Patterns as shown to participants, describing each condition. Based on design from Cauchard et al.'s ActiVibe study.

B. Value perception performance

All trial values were considered missed if Submit was not pressed within 8 s of the trial's vibrations ending. This was generally an adequate amount of time to enter the values, with the Submit button pressed within 6 seconds in the vast majority of trials, and typically faster in Phase 1.

The most basic performance measures are the Miss Rate (MR: % of trials with no value chosen) and Error Rate (ER: % of trials with incorrect response, including MR). The Correct Rate (CR), or 100-ER, is the percentage of trials with the participant submitting exactly the correct answer.

	Miss Rate (%)			Error Rate (%)			Correct Rate (%)		
Cond.	Val1	Val2	Val3	Val1	Val2	Val3	Val1	Val2	Val3
AVF	0.6			24.2			75.8		
BPV	0.3			68.1			31.9		
DO	0.6			74.7			25.3		
AVF3	10.3	10.0	11.4	39.7	44.2	47.5	60.3	55.8	52.5
BPV3	3.1	3.6	1.7	72.5	77.5	73.1	27.5	22.5	26.9
DO3	2.5	1.7	1.9	68.6	70.8	72.8	31.4	29.2	27.2

TABLE II: Miss (MR), Error (ER), and Correct (CR) rates across conditions for Values 1–3.

Table II provides the MR, ER, and CR for each condition. We point in particular to CR, which even in the best performing condition (AVF, single value) is still only 75.8%. For applications where specific number accuracy is crucial (e.g., conveying a phone number), none of the renderings are likely to perform adequately under task demand conditions similar to those in this experiment. Nonetheless, AVF enjoys a clear ER and CR advantage over both BPV and DO. The MR picture is more mixed. In Phase 1, all three rendering conditions have an MR of less than 1%. However, AVF had the largest MR rise (0.6% to 10.3%) when moving from the single to multiple value conditions, with an MR sometimes over triple that of DO3 or BPV3. This may be due to participants simply getting lost with the numerous vibration pulses, such that they are more likely to give up.

These measures do not, however, take into account the magnitude of the errors. To evaluate this, we use the DIA, or absolute value of the difference between the value actually rendered and the perceived value, also used by Cauchard et al. We incorporate the MR into the DIA by assigning missed values a DIA of 10, or one more than the maximum DIA when a value was selected, on the assumption that it is worse to have no idea of a value than to be able to at least make a guess. DIA results are shown in Figure 3.

When comparing DIA results between conditions, we use a non-parametric Friedman test of differences among repeated measures, via the Agricolae package [19] in R [20]. Where the Friedman test showed a significant difference, post-hoc analysis used Agricolae's built-in Friedman LSD pairwise tests, followed by Holm-Bonferroni correction.

For the one value trials (Phase 1), the Friedman test showed a significant difference ($X^2 = 6.94, p = 0.031$) between conditions, with post-hoc pairwise comparison indicating only a difference for AVF-DO (p=0.023). Thus, AVF has better DIA performance than DO, mirroring Cauchard et al.'s results even in the face of the audio distractor task. We cannot claim BPV has better accuracy than DO.

For the three value conditions (Phase 2), there were no significant DIA differences between rendering conditions within each of the three values. However, Friedman tests comparing Phase 1 to 2 DIA, i.e., AVF-AVF3, DO-DO3, and BPV-BPV3, for the first value, showed a significant AVF-AVF3 difference (p=0.039). Thus, AVF had the only statistically significant performance drop, with its mean DIA effectively falling to the level of BPV and DO. AVF-AVF3 deterioration for the 2nd and 3rd values is at least as severe.

C. Audio color name (Blues) task performance

Tapping the Blue button within 3.5 s of a blue stimulus onset was considered acknowledged, and otherwise missed. In pilot testing, the number of missed blue events was highest while the haptic stimuli were being administered, exactly when attention needed to be split between the two tasks. Thus, we focus on the percentage of blue stimuli successfully acknowledged only during the 1 s before through 1 s after the vibrations in each trial. Designs that are shorter to render (e.g., DO), will have fewer overlapping blue stimuli because of their lower duration. However, this is an important consideration, as one would expect that more parsimonious designs would indeed have a practical advantage in terms of distraction vs. renderings that take more time, and therefore demand longer attention. Thus, we do not consider this a critical confound for the current study, since the percentage of missed blue stimuli during the renderings reflects both cognitive load and the effects of rendering duration, both of which are important considerations in real-world use.

Within Phase 1 (AVF, BPV, DO), the mean percentage of acknowledged audio blue stimuli across rendering conditions ranges from 80.1% – 88.7%, and for Phase 2 (AVF3, BPV3, DO3), from 68.7% – 75.7%. Given this performance, we conclude that the Blue task indeed requires attention.

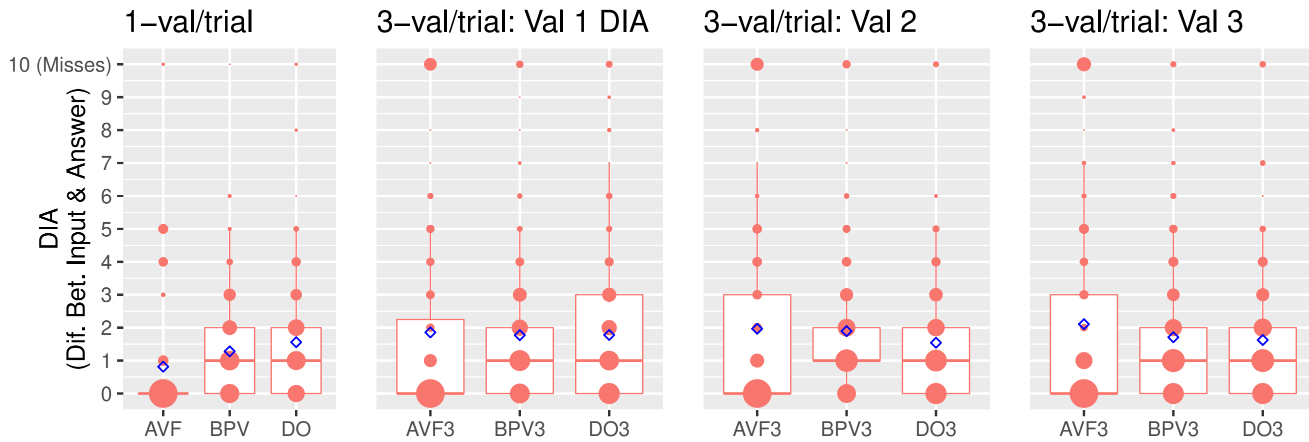

Fig. 3: Accuracy as measured by DIA, with misses assigned DIA=10, by condition. Box plot hinges represent 25% quartiles, whiskers 1.5 * IQR. Blue diamonds represent the mean DIA. Dot sizes are proportional to the number of trials at each DIA level, e.g., the number of misses (DIA 10) for AVF increased between Phase 1 (1-val) and Phase 2 (3-val).

However, even though the mean acknowledgment rate for BPV is slightly higher than for DO or AVF, at less than 10% difference between the conditions within each phase, any performance difference is insignificant for practical purposes. These differences are therefore unlikely to steer application designers between rendering conditions.

D. *Questionnaire results*

Following each phase, participants ranked the three patterns from best (rank 1) to worst (rank 3) for their overall preference, perceived accuracy, and least effort (Table III).

Phase 1: Friedman tests comparing the rankings for each of the three questions showed a significant difference for preference ($X^2 = 9.17, p = 0.02$), with a post-hoc Nemenyi multiple comparison test revealing that AVF is significantly preferred over DO ($p = 0.02$). For perceived accuracy ($X^2 = 13.5, p = 0.001$), participants' subjective evaluations were again clearly in favor of AVF over DO ($p = 0.0007$).

Phase 2: None of the subjective measures had a statistically significant difference between rendering conditions. Thus we conclude that as with DIA, AVF enjoys a statistically significant advantage over DO in Phase 1, but this disappears in the more difficult three-value condition. Further, AVF does not have a statistically significant subjective advantage over BPV in any of the three questions in both phases, but this is particularly evident in Phase 2, where for effort and overall preference, BPV ties for or receives the (again, not statistically significant) best mean rank.

Some participants also wrote additional comments. For example, p02 reported their strategy for the three approaches as, "AVF – count, BPV – compare, DO – guess," indicating they understood the fundamental differences between the rendering options. Separately, p11 indicated that the "reference vibration in BPV made it easier to distinguish," and in reference to DO, "it comes out of nowhere" suggesting that they perceived at least some value in the reference vibration.

Phase	Question	DO	BPV	AVF
1 (1-val)	Accuracy	*2.75	2.00	***1.25**
	Least effort	2.00	2.33	**1.67**
	Preference	*2.42	2.25	***1.33**
2 (3-val)	Accuracy	2.33	1.92	**1.75**
	Least Effort	**1.75**	**1.75**	2.50
	Preference	1.92	**1.83**	2.25

TABLE III: Mean rankings of conditions. Preferred option has rank 1. Bold font indicates the best ranked choice, * represents significant ($p < 0.05$) difference within row.

VI. Discussion

In Phase 2, DO3 and BPV3 perform about as well as AVF3 both in mean DIA and in subjective preference. If a low MR is crucial to an application, but a higher ER is tolerable, then either DO or BPV may be preferable in conditions similar to Phase 2, since AVF suffers a substantial increase in MR vs. Phase 1. When taking into consideration ActiVibe's greater total time to render values, as well as associated increased power consumption from running the vibration motor longer (Table I), DO may be a legitimate option for applications that render frequent multiple values in the face of ongoing distraction, such as the SenseProxy application (Section I).

However, the hope that BPV, with its perceptual calibration PreVibe, would prove superior to AVF is not supported. Participants remarked on the difficulty of the overall task, e.g., "I found the 3 vibrations a lot harder than 1—much harder to keep track & remember them while listening to the colours, & hard to distinguish between the three vibrations." (p11) Thus, we conjecture that any benefit from the BPV perceptual calibration may be undermined by adding a fourth vibration that needs to be processed, and can be confused with the information vibrations. In the end, BPV appears practically no better than DO.

VII. Limitations and Future Work

Our results are limited to vibrotactile patterns on the wrist. Generalizing to other body locations may impact performance, and using non-vibrotactile actuators will become increasingly pertinent as commercial products incorporate a wider variety of haptic hardware. Testing alternative distraction conditions and a broader range of information content (e.g., 2 or 4 values) may better clarify tradeoffs between pattern complexity, accuracy, and subjective preference.

More training time may have greater benefits for some renderings, which would potentially change the cost/benefit balance between rendering methods. An alternative training regimen, e.g., Passive Haptic Learning [21], may also help improve performance and offer an efficient way to help users learn to interpret the haptic patterns.

Vibration parameters including the length of the gap between multiple values and duration increments were validated with limited pilot testing. AVF, in particular, may benefit from longer gaps between values, although this would increase rendering time. Nonetheless, further optimizing the patterns may improve results for all three approaches. Using vibrotactile actuators that allow changes in parameters such as frequency may also allow for improved pattern designs.

Confounds such as motion during a haptic stimulus can cause a stimulus to be perceived as less intense [22], or missed entirely [23]. Indeed, ActiVibe's accuracy fell from an overall 88.7% accuracy to 54.1% while the participant was running [2]. An untested hypothesis, since participants were at rest during this experiment, is that the BPV calibration vibration would significantly help interpret following vibrations specifically when perception is warped by motion.

The training and trials occurred in a short period of time. This likely improves DO's performance vs. BPV, since remembering the durations from training is easier than if values were administered sporadically throughout the day. Indeed, ActiVibe participants sometimes confused even the short (150 ms) and long (600 ms) vibrations [2]. Thus, BPV may show an advantage over DO if the reminder PreVibe proves more valuable as memory of DO durations fades.

Last, to be able to directly compare to ActiVibe's results, we rendered values from 1–10. However, based on our DIA error results, for applications requiring fewer gradations (e.g., low/medium/high), we expect that DO and BPV are sufficient, and would likely be preferred to AVF, which would be overkill for distinguishing only a few levels.

VIII. Conclusion

In sum, our results indicate that although ActiVibe indeed has the best ER and CR across the board, its single value DIA, MR, and subjective preference advantages can be largely eroded when rendering multiple values in a high-distraction environment, providing insight into the limits of vibrotactile numeric information delivery.

Acknowledgment

The authors wish to thank Pascal Fortin and the other members of the Shared Reality Lab for feedback and testing.

References

[1] J. R. Blum and J. R. Cooperstock, "Exploring the Limits of Vibrotactile Numeric Information Delivery," in *Proc. Intl. Symposium on Pervasive and Ubiquitous Computing and Wearable Computers - UbiComp '18*. New York, USA: ACM Press, 2018, pp. 17–20.

[2] J. R. Cauchard, J. L. Cheng, T. Pietrzak, and J. A. Landay, "ActiVibe: Design and Evaluation of Vibrations for Progress Monitoring," in *Proc. 2016 CHI Conf. on Human Factors in Computing Systems - CHI '16*. New York, USA: ACM Press, 2016, pp. 3261–3271.

[3] S. Brewster and L. Brown, "Tactons: structured tactile messages for non-visual information display," *Proc of the Fifth Conference on Australasian User Interface*, vol. 28, no. January, pp. 15–23, 2004.

[4] J. R. Blum and J. R. Cooperstock, "Expressing Human State via Parameterized Haptic Feedback for Mobile Remote Implicit Communication," in *Proc. 7th Augmented Human International Conference - AH '16*. New York, USA: ACM Press, Feb. 2016, pp. 1–2.

[5] K. MacLean, "Putting Haptics into the Ambience," *IEEE Transactions on Haptics*, vol. 2, no. 3, pp. 123–135, July 2009.

[6] S. Brewster and A. King, "An Investigation into the Use of Tactons to Present Progress Information," in *Lecture Notes in Computer Science*, 2005, vol. 3585 LNCS, pp. 6–17.

[7] M. Pielot and R. de Oliveira, "Peripheral Vibro-Tactile Displays," *Proc. 15th Intl. Conf. on Human-Computer Interaction with Mobile Devices and Services (MobileHCI'13)*, pp. 1–10, 2013.

[8] I. Oakley and J. Park, "Did you feel something? Distracter tasks and the recognition of vibrotactile cues," *Interacting with Computers*, vol. 20, no. 3, pp. 354–363, 2008.

[9] B. Saket, C. Prasojo, Y. Huang, and S. Zhao, "Designing an effective vibration-based notification interface for mobile phones," in *Proc. Conf. on Computer Supported Cooperative Work - CSCW '13*. New York, USA: ACM Press, 2013, p. 149.

[10] L. M. Brown, S. A. Brewster, and H. C. Purchase, "Multidimensional Tactons for Non-Visual Information Presentation in Mobile Devices," *Proc. 8th Conf. on Human-computer interaction with mobile devices and services - MobileHCI '06*, vol. 159, no. Sep., pp. 231–238, 2006.

[11] A. Tang, P. McLachlan, K. Lowe, C. R. Saka, and K. MacLean, "Perceiving ordinal data haptically under workload," in *Proc. 7th Intl. Conf. on Multimodal interfaces - ICMI '05*. New York, New York, USA: ACM Press, 2005, p. 317.

[12] Y. Wang, B. Millet, and J. L. Smith, "Designing wearable vibrotactile notifications for information communication," *International Journal of Human Computer Studies*, vol. 89, pp. 24–34, 2016.

[13] E. C. Lechelt, "Temporal Numerosity Discrimination: Intermodal Comparisons Revisited," *British Journal of Psychology*, vol. 66, no. 1, pp. 101–108, Feb. 1975.

[14] A. Bianchi, I. Oakley, and D. S. Kwon, "Counting clicks and beeps: Exploring numerosity based haptic and audio PIN entry," *Interacting with Computers*, vol. 24, no. 5, pp. 409–422, 2012.

[15] T. Philippi, J. van Erp, and P. Werkhoven, "Multisensory temporal numerosity judgment," *Brain Research*, vol. 1242, pp. 116–125, 2008.

[16] J. Pasquero, S. J. Stobbe, and N. Stonehouse, "A haptic wristwatch for eyes-free interactions," in *Proc. Conf. on Human factors in computing systems - CHI '11*. New York, USA: ACM Press, 2011, p. 3257.

[17] M. C. Frank, E. Fedorenko, P. Lai, R. Saxe, and E. Gibson, "Verbal interference suppresses exact numerical representation," *Cognitive Psychology*, vol. 64, no. 1, pp. 74–92, 2012.

[18] A. Chan, K. MacLean, and J. McGrenere, "Learning and Identifying Haptic Icons under Workload," in *1st Eurohaptics Conf. and Symp. on Haptic Interfaces for Virtual Environment and Teleoperator Systems*. IEEE, 2005, pp. 432–439.

[19] F. de Mendiburu, *agricolae: Statistical Procedures for Agricultural Research*, 2017.

[20] R Core Team, "R: A Language and Environment for Statistical Computing," Austria, 2014.

[21] C. Seim, S. Reynolds-Haertle, S. Srinivas, and T. Starner, "Tactile taps teach rhythmic text entry: passive haptic learning of morse code," *Procp. 2016 ACM Intl. Symposium on Wearable Computers - ISWC '16*, pp. 164–171, 2016.

[22] A. Morrison, L. Knudsen, and H. J. Andersen, "Urban Vibrations: Sensitivities in the Field with a Broad Demographic," in *Proc. Intl. Symp. on Wearable Computers*, vol. 1. IEEE, June 2012, pp. 76–79.

[23] J. R. Blum, I. Frissen, and J. R. Cooperstock, "Improving Haptic Feedback on Wearable Devices through Accelerometer Measurements," *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology - UIST '15*, pp. 31–36, 2015.