Improving Haptic Feedback on Wearable Devices through Accelerometer Measurements

Jeffrey R. Blum[1,2] **and Ilja Frissen**[1,3] **and Jeremy R. Cooperstock**[1,2]
McGill University, Montreal, Canada
[1]Centre for Interdisciplinary Research in Music Media and Technology (CIRMMT)
[2]Shared Reality Lab, Centre for Intelligent Machines; [3]School of Information Studies
[jeffbl,jer]@cim.mcgill.ca, ilja.frissen@mcgill.ca

ABSTRACT

Many variables have been shown to impact whether a vibration stimulus will be perceived. We present a user study that takes into account not only previously investigated predictors such as vibration intensity and duration along with the age of the person receiving the stimulus, but also the amount of motion, as measured by an accelerometer, at the site of vibration immediately preceding the stimulus. This is a more specific measure than in previous studies showing an effect on perception due to gross conditions such as walking. We show that a logistic regression model including prior acceleration is significantly better at predicting vibration perception than a model including only vibration intensity, duration and participant age. In addition to the overall regression, we discuss individual participant differences and measures of classification performance for real-world applications. Our expectation is that haptic interface designers will be able to use such results to design better vibrations that are perceivable under the user's current activity conditions, without being annoyingly loud or jarring, eventually approaching "perceptually equivalent" feedback independent of motion.

Author Keywords

Wearable computing; mobile sensing; haptic vibration feedback; accelerometer

ACM Classification Keywords

H.5.2 : User Interfaces—*Haptic I/O*

INTRODUCTION

Haptic feedback on mobile devices is intended to provide subtle notifications while on-the-go, and is thus an increasingly important part of mobile, and especially wearable, devices. We observe that it is easy to notice a smartwatch vibration while sitting in a conference room, yet while riding a bike on a bumpy path, the same stimulus can go unnoticed. Indeed, as discussed in Related Work, coarse activity measures such as moving from a laboratory environment to outdoors or classifying whether the person is walking, impacts perception. Thus, we hypothesized that motion, as measured by an accelerometer just before a vibration stimulus is delivered, may offer a significant improvement in determining whether the stimulus will be perceived. This paper describes an experiment that uses a smartwatch to measure motion just before administering different vibration stimuli throughout the day, and records whether the user notices them. We only used sensors built into current smartwatches, eschewing external sensors and mechanisms. Although this limits the sensing capabilities, the results can be deployed immediately for practical notification and haptic communication applications.

If prior acceleration is indeed predictive, it would deliver additional benefits beyond the main feature of more subtle, yet still effective vibration feedback regardless of activity level. First, since the MEMS accelerometer in the Pebble smartwatch (ST LIS3DH) consumes about .036 mW, while even a small eccentric rotating mass (ERM) vibration motor (e.g., Precision Microdrives 310-004) consumes about 42 mW, we expect that using the accelerometer to reduce vibration duration or intensity will also provide a net power savings. In addition, it would also enable the device to wait a brief period until there is less motion before administering a weaker stimulus, further reducing power consumption and potentially allowing the use of smaller, less expensive, or less power-intensive actuators. If motion does not drop below a level where the stimulus would likely be felt, even with the actuator driven at maximum intensity, the system could fall back to audible or other methods to ensure an important notification is perceived. Although such tradeoffs are not necessarily as critical in larger products such as tablets or smartphones, wearable devices must aggressively minimize size and power consumption in order to be comfortable, stylish, and usable.

RELATED WORK

Creating noticeable but not overwhelming vibrations has been a long-standing topic of interest. Researchers at Nokia studied how vibration duration changes subjective perception of the stimulus, finding that durations between 50-200 ms are perceptible from a phone in a front trouser pocket in a laboratory setting, with ambient sound blocked by headphones playing pink noise [6]. At 500 ms duration, vibrations were reported as "too strong" much more frequently than at 200 ms.

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. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
UIST 2015, November 8–11, 2015, Charlotte, NC, USA.
Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-3779-3/15/11 ...$15.00.
DOI: http://dx.doi.org/10.1145/2807442.2807474

Studies in human and animal perception indicate that active movement can lead to a partial suppression, or "gating", of the transmission of tactile inputs [12]. For example, in contrast to the stationary participants in the Nokia study, Martinsson reported that increasing the duration of single pulse vibrations showed clear perception benefits up to 800 ms while walking [9], and Qian et al. likewise found that 800 ms duration haptic icons performed better for ambulatory users than those at 200 ms [13]. Baek et al. found that the optimal frequency of vibration for best perception was higher while walking rather than when stationary [2]. Karuei et al. studied vibrations of 500 ms duration at 12 points on the body, including the wrist, and found that, "Walking significantly reduces odds of detecting a vibration," but that visual workload had no, "apparent effect on vibration detection" [7]. Additional variables that impact sensitivity to a vibration stimulus include amplitude [3], sound [13] and temperature [5].

Similar to the method presented here, Pasquero et al. had participants note when they felt vibrations, rendered by a 30 mm radius piezoelectric actuator in a custom-built watch, during their normal activities over a two-hour period in a workplace setting [11]. Although the participants achieved very high detection rates (97%, with durations ranging from 256 to 1280 ms), it is unclear how well these results apply to more standard vibration actuators, and in contexts outside of a workplace. The study did not measure device motion.

A comparison of vibration perception via a haptic belt, with participants splitting their time between a laboratory and walking in an urban environment, was conducted by Morrison et al. [10]. They found a, "lowering of sensitivity in the field (with distraction)." Although conducted in an urban setting, this study was short-term (both lab and field conditions in less than one hour per participant), and did not measure motion specifically before each vibration.

Andersen et al. tested various predictors of vibration perception and showed that logistic regression is a good model for such a perception study. They found that age, stimulus intensity and "situation" were significant predictors, with "situation" showing a decrease in perception when moving from a laboratory environment to an outdoor setting, and older participants demonstrating a greater drop in perception when switching to the outdoor context. They conclude as follows:

> "for real life implementation the situation cannot be so pre-determined but has to be estimated semi or fully automatically to adapt vibration intensity to the optimal level. . . various sensors monitoring biological and motion information. . . could be sampled with fusion to give an estimate of the current situation." [1]

The work presented in this paper does exactly this using a sensor already present in a commercial smartwatch to gauge the user's current state not at a gross level, but precisely at the location of stimulus and just before the vibration is triggered. We also use the same form of analysis, logistic regression, to extend their results in the direction they suggest.

In summary, this paper presents a considerable practical improvement over the cited work, including that of Andersen et al., whose coarse activity "situation" was not measured by a device. Even dynamically categorizing context using sensors (e.g., "walking") would likely miss brief lulls in which a stimulus would be perceptible. Conversely, a resting person's transient wrist motion could mask a stimulus presented at just the wrong time. In contrast, our primary contribution is showing that prior acceleration, a physically and temporally proximate, easy-to-measure factor, significantly improves predictions of whether a stimulus will be perceived. In addition to the less jarring notifications and potential power savings already described, this work advances the field as, to our knowledge, the first demonstration of a technique that is practically implementable on an inexpensive device without explicit categorization of the user's activity, and is much more responsive to transient changes in activity than classification approaches, requiring sensor data less than 1 s before the stimulus to make a decision. Further, the "in the wild" nature of the study illustrates that the effect is evident despite confounding variables such as ambient noise and participant distraction.

METHOD

We recruited seven volunteer participants (5 male, 2 female, ages 19-47, median=24), each of whom wore an original model Pebble smartwatch, containing a pancake ERM vibration motor, for at least three days, not necessarily consecutively. The watch was placed on the bare wrist where the participant normally wore a watch. For the 4 participants who did not normally wear a watch, it was placed on the wrist where they indicated they would prefer to wear one, which was the non-dominant hand in all cases. The strap was made only as tight as was comfortable, with the experimenter verifying that it was not so loose as to be uncoupled, and participants were instructed to always use the same clasp position.

When the participant noticed a vibration, they were to press the lower-right button on the smartwatch to indicate they had perceived the stimulus. For safety reasons, participants were instructed *not* to push the button immediately after feeling a stimulus if it was not safe to do so. In such a case, they could instead push the upper-right button to indicate they had perceived the vibration but were unable to press the button immediately. If the timestamp of a vibration was followed within 8 s by a lower button press, or within 60 s by an upper ("late") button press, the stimulus was considered perceived, and otherwise missed. In practice, the "late" button was only used six times during the entire experiment to indicate a perceived vibration after 8 s. Participants were typically quick to push the lower right button (mean delay=2.8 s, sd=1.8 s). The Pebble application displayed text beside each button as a reminder of which button to push. When a button was pushed, a message appeared briefly in the middle of the watch's screen as confirmation (Figure 1).

The watch vibrated with different stimuli every three to ten minutes. Although there is prior work evaluating variables such as vibration intensity and pause duration in patterns, we focused on only a single pulse so as to determine whether there is indeed a predictive effect in the simplest case. A pilot experiment with eight participants found that stimuli at or above 300 ms duration and 4/10 of maximum intensity were

(a) Pebble watch

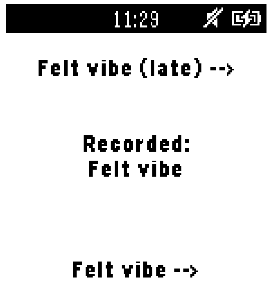

(b) UI just after pressing lower-right side button

Figure 1: Pebble smartwatch.

perceived the vast majority of the time regardless of prior acceleration, so we chose smaller values for this experiment. Intensity was varied via pulse width modulation (PWM) pulsing with a 10 ms duty cycle[1] at values of either 2 (vibrating 2 ms out of every 10 ms) or 4. Vibration duration was either 100 ms or 200 ms, making four total vibration duration/strength conditions, referred to as 100/2, 200/2, 100/4, 200/4 throughout this paper. The four combinations were administered in random order, then re-randomized for subsequent stimuli.

Prior to each stimulus, three-axis accelerometer readings were acquired at 50 Hz, with a range of ±4 g. The accelerometer data was filtered using a basic high-pass filter with a 0.1 s time constant[2] to remove gravity and keep only dynamic acceleration. We note there is research into how best to filter and use accelerometer data to measure physical activity levels [17], but since we are examining the more specific issue of motion relevant to vibration perception, it is unclear how relevant such results are to our use case. The magnitude of the three-axis dynamic acceleration vector was averaged over all the samples recorded in the 500 ms before the stimulus was initiated, which we refer to as "prior acceleration". In an earlier pilot experiment, we found that most acceleration values were very low since participants were largely sedentary. We considered using activities such as sports, but wanted the results to reflect mostly normal, everyday motions. Instead, the Pebble application monitored the current acceleration level in order to bias trials toward higher acceleration values, as follows. Each time a stimulus was scheduled to begin, the application started to look for a mean dynamic acceleration value above 70 milligs over the prior 0.1 s, and dropped that threshold to zero linearly over a 30 s period, triggering a vibration when the threshold was exceeded. Thus, a vibration would always trigger by the end of the 30 s window, but the application would "catch" higher activity levels to get a better distribution of prior acceleration values. In order to avoid the participant learning that a sudden motion would trigger a vibration due to this algorithm, it immediately vibrated in 25% of trials, regardless of the initial acceleration level.

[1]Code available at: https://github.com/jeffbl/pebble
[2]Based on https://developer.android.com/guide/topics/sensors/sensors_motion.html

Because the watch can only store a limited amount of data in memory, each participant also carried an Android smartphone or tablet that received and stored the data until it could be downloaded. Some data was lost due to issues with the Pebble data logging system, which appeared to be triggered when the watch was too far from the logging device for an extended period. In these cases, this required extending the time with the device to gather additional data. One participant reported falling asleep while the application was running, so this time period (approx. 8pm - 7am) was excluded from analysis. In the end, we obtained at least 50 vibrations per participant per condition, and a total of 3221 vibration stimuli across all participants and conditions. Specific counts are summarized in Table 1. Most notably, p001 ran for a longer period than the other participants, and thus accounts for roughly a third of the total data points. At the other extreme, p003 has the fewest data points, largely due to the data logging issue.

RESULTS

As advocated by Andersen et al. [1], we use logistic regression, a statistical method for analyzing binary outcomes based on one or more predictor variables, to estimate the probability of a stimulus being perceived. Data analysis was via R [14], using its glm function with family=binomial. In this section, we first show a link between prior acceleration and perception on an individual level. We then provide quantitative results of a logistic regression model incorporating pooled data across all participants. Last, we measure the performance of a classifier based on these logistic regression results.

Logistic regression results

To get a better feel for the data we first explore a subset of the results in more detail. For the benefit of clarity, we focus initially on a single participant p001 in the 100/4 condition, shown in Figure 2a. The dots in the plot represent the percentage of stimuli perceived when the data is binned in 25 millig prior acceleration increments, purely for visualization, since plotting individual stimuli results in an unintelligible smear at the binomial 0 (missed) and 1 (perceived) probability levels.

First, we note that the data is heavily biased toward lower prior acceleration values (larger dots, representing more vibrations pooled into the 25 millig bin), since participants were not always active. We have a robust number of samples up to approximately 200 milligs, after which it rapidly becomes sparse, meaning that we are less sure of the perception rates at higher prior accelerations. As we will see, perception falls off at less than 200 milligs of prior acceleration, especially for the lower stimulus levels. Next, we note that the maximum perception probability occurs at lower prior acceleration values, i.e., when the wrist is nearly motionless, with perception falling off as prior acceleration increases.

Figure 2b shows individual logistic regression curves at the weakest stimulus level (100/2). In contrast to the 100/4 condition shown in Figure 2a, we see a much faster falloff in perception for p001 (now orange), due to the weaker stimulus. Perception rates for some participants do not even rise to 75%, indicating the stimulus was not always perceptible even with minimal wrist motion. Note individual differences,

		duration (ms)/intensity							
		100/2		200/2		100/4		200/4	
	age	vibes	%felt	vibes	%felt	vibes	%felt	vibes	%felt
p004	19	70	57%	74	70%	71	70%	70	87%
p007	19	104	55%	105	75%	105	70%	105	90%
p005	21	113	34%	110	83%	113	79%	110	97%
p001	24	255	29%	255	71%	253	68%	252	91%
p006	25	90	28%	88	40%	89	55%	88	74%
p003	30	51	41%	52	85%	50	78%	53	92%
p002	47	123	15%	123	38%	125	36%	124	69%

Table 1: Summary per participant, sorted by age: total vibrations and percentage perceived for each duration/intensity condition.

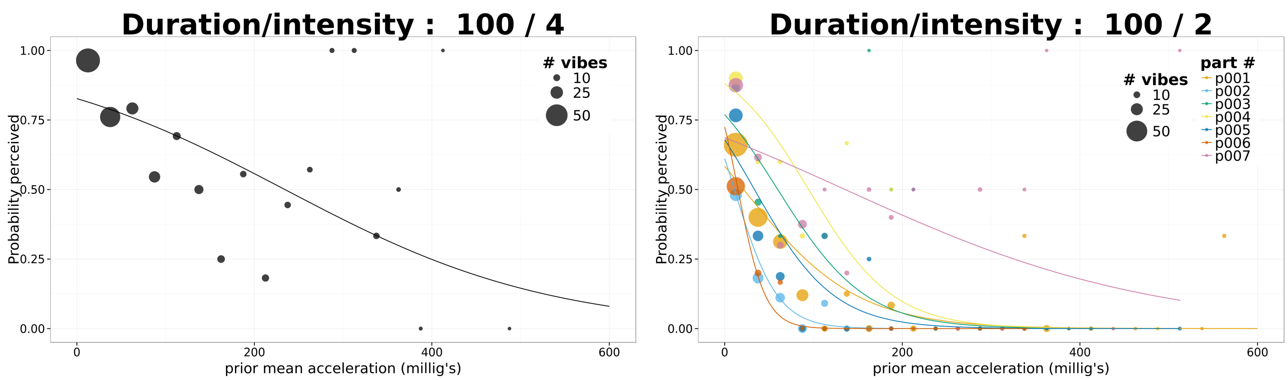

(a) Participant p001, 100/4 condition.

(b) Regression per participant for 100/2 condition.

Figure 2: Perception data and regressions on a per-participant basis.

such as p007 being substantially more sensitive at higher prior acceleration levels. At the other extreme, p006 and p002 show a faster decline in perception. Except for p007, only a small proportion of vibrations are perceived above a 200 millig threshold. Note that although there is significant inter-participant variation when looking only at duration and intensity as predictors, the overall regression results take into account participant age, which is not apparent from Figure 2b.

Although not shown here, the plot for the 200/4 condition shows much more limited value to the prediction simply because a much higher percentage (86%) of the vibrations are perceived overall. Above this stimulus strength, we expect based on an earlier pilot test that practically all vibrations will be perceived, at least in the observed acceleration ranges.

We now build an overall logistic regression model using pooled data from all participants and conditions, including four predictor variables: intensity, duration, age, and prior acceleration. Table 2 summarizes the logistic regression results. As expected based on the Related Work above, intensity and duration are significant predictors, as demonstrated through the small p-values computed from a χ^2 likelihood ratio test (marked L.R.T. $p > \chi^2$ in Table 2), generated using R's anova function with test=Chisq, which is equivalent to test=LRT when performing logistic regression. Since Andersen et al. showed that age is also important, we included it in our model, and confirmed that it is indeed a significant predictor, although it is important to note that our participants were in a more restricted age range (19-47 vs. 7-79 years) [1]. Most importantly for the purposes of the present study, we found that prior acceleration is also a significant predictor.

A separate test was performed to evaluate the final residual deviance of the overall model, which is a measure of the overall model fit to the data. Even if this final fit turned out to be poor, we would conclude, based on the above analysis, that prior acceleration does significantly improve the prediction, but that there are likely unaccounted for variables that would be necessary to achieve a good overall fit to the observed data. For this test, as explained by Cook et al., "...the null hypothesis states that the logistic regression model provides an adequate fit to the data." [4]. Since the chance that a χ^2 with 3216 degrees of freedom exceeds the final residual deviance of 3259 is 0.294 ($p > 0.05$), we *cannot* reject this null hypothesis, and thus conclude that not only is prior acceleration an important contribution to the model, but also that the model containing all four predictors has a reasonably good fit to the data. This residual deviance should not be surprising given that the experiment was carried out in the field, with uncontrolled variables such as ambient sound levels and participant distraction, both of which are likely sources of variation.

Classification performance

A practical application would likely use the logistic regression model to classify whether a vibration with given parame-

	Coeff (β)	Std.Err	Residual Deviance	Residual Df.	Deviance Reduction	L.R.T. $p > \chi^2$
Intercept (β_0)	-1.4625	0.2139	4263	3220		
duration	0.0154	0.0009	4012	3219	251	< 0.0001
intensity	0.7386	0.0454	3766	3218	246	< 0.0001
age	-0.0616	0.0047	3602	3217	164	< 0.0001
prior accel	-0.0072	0.0004	3259	3216	344	< 0.0001

Table 2: Logistic regression results for all data, showing prior acceleration as a predictor significantly reduces residual deviance (χ^2 Likelihood Ratio Test $p < 0.05$), as do vibration intensity/duration and age.

ters would be perceived. We evaluate this with 10-fold cross-validation (via the caret package [8]), building mean ROC curves using ROCR [16], based on the classification performance of the ten resulting logistic regression models.

Figure 3 shows the mean ROC curves (vertically averaged) for a model containing only duration+intensity+age, then adding prior acceleration. The mean area under the curve (AUC) for the first model is 0.757, while that of the model containing prior acceleration rises to 0.817, indicating that the latter is indeed superior. DeLong's test, via the pROC library [15], on combined ROC curves from the 10-fold cross-validation indicates $p < .0001$, so we reject the hypothesis that the true difference in AUC is zero. Starting at the left of the plot, we see that the performance of both models tracks closely together until reaching a true positive rate (TPR) of around 0.30. After this point, the model performance diverges, with the model including prior acceleration maintaining a lower false positive rate as its TPR climbs.

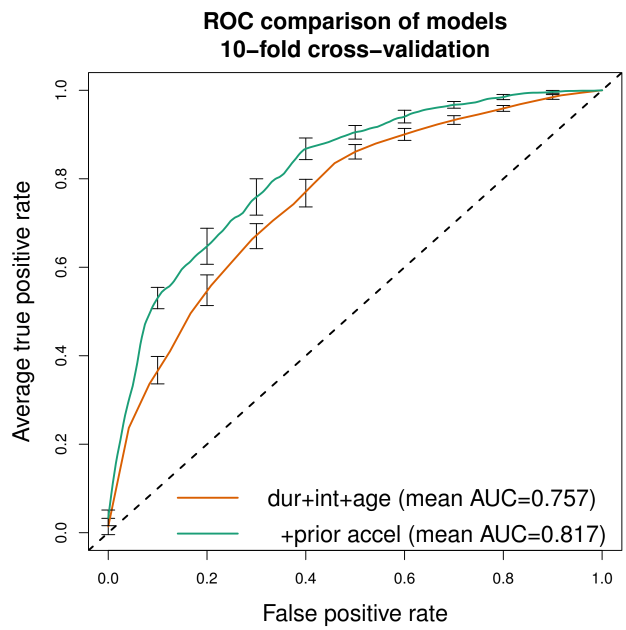

Figure 3: ROC curve comparison of classification performance for models without ("dur+int+age") and with ("+prior accel") prior acceleration as a predictor. Error bars represent 1.96 × the standard error of the models generated by 10-fold cross-validation, providing a 95% confidence interval.

As noted earlier, there was a considerable difference in the number of data points gathered for each participant. Since we were concerned that the results may be overly optimistic due to, for example, approximately a third of the data coming from p001 alone, we carried out two additional 10-fold cross-validation runs. The first upsampled the data for each participant to the same total number of samples as p001, while the second downsampled each participant to the same total number as p003. The mean AUC values across the 10 folds were not much different from those generated using the unbalanced data set, as shown in Table 3. Most importantly, the difference between the AUC values ("Difference" in the table) for the two models is roughly consistent.

	Mean AUC		
	dur+int+age	+prior accel	Difference
unbalanced	0.757	0.817	0.060
upsampled	0.740	0.811	0.071
downsampled	0.728	0.795	0.067

Table 3: Mean AUC ROC results for 10-fold cross-validation on unbalanced (pooled), upsampled and downsampled data.

FUTURE WORK

First, we would like to extend our work to encompass ambient sound levels, which also impact whether vibrations are successfully perceived [13]. Further improvements in the analysis of data already acquired should also be pursued. For example, averaging of overall acceleration magnitudes across all three axes may mask effects related to the direction of motion, e.g., motion aligned with gravity, or in/out of alignment with the plane of the vibration motor may be important considerations. Likewise, some of the parameters chosen in our study were selected as "educated guesses" but were not optimized, including the choice of averaging 500 ms of acceleration prior to each vibration as a predictor. Ideas include changing the amount of time, using an averaging method biased toward later (more proximate to the stimulus) acceleration samples, and using squared acceleration values to give more weight to large acceleration impulses. Other factors, such as the tightness of the watch coupling to the wrist, could potentially be measured by the accelerometer since one would expect the vibration to be more damped as band tightness increases, providing a rough measure of the coupling.

Other regressors or machine learning techniques could potentially improve our classification results, but we believe these will be best pursued when moving to a more inclusive model encompassing more devices and data from a broader participant pool, and when targeting a specific application.

Despite efforts to bias the data collection to periods where the participant had greater wrist acceleration prior to the stimulus, the data was still heavily skewed in favour of low acceleration values. We expect that further efforts to gather data at these higher acceleration levels would refine the curves.

CONCLUSION

The duration and intensity of a vibration stimulus given by a smartwatch can be insufficient to reliably predict whether a vibration will be perceived. We hypothesized that measuring the amount of watch motion via the accelerometer before the onset of vibration stimulus would prove useful in making better predictions. Through our experiment, we found a significant predictive correlation between prior acceleration of the wrist and whether a vibration stimulus is perceived across seven participants. This effect was very clear and significant at the lower stimulus levels, becoming less important when a participant perceived the stimulus most of the time regardless of prior acceleration. Such predictions are useful in practice since our algorithm is easily implementable on even low-end wearable devices available today, and would allow more restrained and potentially more power efficient haptic feedback to be used, rather than the "one size fits all" approach prevalent today. We hope that these results can be refined for prior acceleration, as well as extended to other variables such as sound and attention, to create "perceptually constant" haptic stimuli, opening new possibilities in haptic interfaces.

ACKNOWLEDGEMENTS

The authors thank the participants from McGill's Shared Reality Lab, as well as Vincent Hayward for discussion. This work was funded by an NSERC PGS-D Scholarship, a McGill Engineering Doctoral Award (MEDA), a McGill Lorne Trottier Engineering Graduate Fellowship and a Graphics, Animation and New Media (GRAND) Scholar stipend. McGill Review Ethics Board file #83-0814.

REFERENCES

1. Andersen, H. J., Morrison, A., and Knudsen, L. Modeling vibrotactile detection by logistic regression. In *Proc 7th Nordic Conference on Human-Computer Interaction - NordiCHI '12*, ACM Press (New York, New York, USA, 2012), 500.
2. Baek, Y., and Myung, R. Have you ever missed a call while moving? : The Optimal Vibration Frequency for Perception in Mobile Environments. In *Proc 6th WSEAS Intl Conf on Applied Informatics and Communications*, vol. 2006 (2006), 241–245.
3. Bochereau, S., Terekhov, A., and Hayward, V. Amplitude and duration interdependence in the perceived intensity of complex tactile signals. In *Haptics: Neuroscience, Devices, Modeling, and Applications*, M. Auvray and C. Duriez, Eds., LNCS. Springer Berlin Heidelberg, 2014, 93–100.
4. Cook, D., Dixon, P., Duckworth, W. M., Kaiser, M. S., Koehler, K., Meeker, W. Q., and Stephenson, W. Beyond Traditional Statistical Methods. Iowa State University Press, 2000, ch. 3.
5. Green, B. The effect of skin temperature on vibrotactile sensitivity. *Perception&Psychophysics 21*, 3 (May 1977), 243–248.
6. Kaaresoja, T., and Linjama, J. Perception of Short Tactile Pulses Generated by a Vibration Motor in a Mobile Phone. In *First Joint Eurohaptics Conf and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, IEEE (2005), 471–472.
7. Karuei, I., MacLean, K. E., Foley-Fisher, Z., MacKenzie, R., Koch, S., and El-Zohairy, M. Detecting vibrations across the body in mobile contexts. In *Proc 2011 conf Human factors in computing systems - CHI '11*, ACM Press (New York, NY, USA, 2011), 3267.
8. Kuhn, M. Building predictive models in R using the caret package. *J Stat Softw 28*, 5 (11 2008), 1–26.
9. Martinsson, M. *Detection of mobile phone vibrations during walking*. M. thesis, U of Lund, Sweden, 2011.
10. Morrison, A., Knudsen, L., and Andersen, H. J. Urban vibrations: Sensitivities in the field with a broad demographic. In *Proc Intl Symp on Wearable Computers, ISWC*, IEEE (2012), 76–79.
11. Pasquero, J., Stobbe, S. J., and Stonehouse, N. A haptic wristwatch for eyes-free interactions. In *Proc 2011 Conf on Human factors in computing systems - CHI '11*, ACM Press (New York, USA, 2011), 3257–3266.
12. Post, L. J., Zompa, I. C., and Chapman, C. E. Perception of vibrotactile stimuli during motor activity in human subjects. *Exp Brain Res 100*, 1 (Jan. 1994), 107–20.
13. Qian, H., Kuber, R., and Sears, A. Towards developing perceivable tactile feedback for mobile devices. *Intl J Human-Computer Studies 69*, 11 (Oct. 2011), 705–719.
14. R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Austria, 2014.
15. Robin, X., Turck, N., Hainard, A., Tiberti, N., Lisacek, F., Sanchez, J.-C., and Müller, M. pROC: an open-source package for R and S+ to analyze and compare ROC curves. *BMC Bioinformatics 12* (2011), 77.
16. Sing, T., Sander, O., Beerenwinkel, N., and Lengauer, T. ROCR: visualizing classifier performance in R. *Bioinformatics 21*, 20 (2005), 7881.
17. van Hees, V. T., Gorzelniak, L., Dean León, E. C., Eder, M., Pias, M., Taherian, S., Ekelund, U., Renström, F., Franks, P. W., Horsch, A., and Brage, S. Separating movement and gravity components in an acceleration signal and implications for the assessment of human daily physical activity. *PloS one 8*, 4 (Jan. 2013).