Evaluation of Haptic Feedback in Android Widgets

Project group "Hippity Hoppity Haptical"

Federico Landorno
12/12/1994
landorno@kth.se

Alessandro Iop
16/09/1996
aiop@kth.se

ABSTRACT

UPDATED—July 9, 2020. Despite the high amount of Android applications being developed every year, a set of guidelines regarding how to implement haptic feedback on some of the most common widgets does not exist. In this report different vibration patterns, associated to multiple versions of six selected Android widgets, are evaluated through a mobile application and an associated online survey. Results from 22 collected responses show how "natural-ness", effectiveness and attention grabbed are affected by different parameters such as the number and length of vibrations. Finally, several tips to consider when designing and developing the user interface of an Android application are presented to the reader.

Author Keywords

Haptic Feedback; Android Smartphones; Widget; Vibration; Haptic Interaction; User Interface Development.

CCS Concepts

•**Human-centered computing** → *Smartphones;* User interface toolkits; *User interface design;* •**Hardware** → Touch screens;

INTRODUCTION

As often happens during the design (and later development) process of applications for mobile platforms, output modalities such as the visual and auditory ones are the most focused on with regards to interaction with the user interface (UI). Haptic feedback, which consists primarily of vibrations conveyed through a touch screen, is meanwhile underestimated when it comes to provide a quality user experience. In a domain of application pertaining entertainment such as gaming, it was however verified by Choe and Schumacher [3] that the inclusion of vibration *"significantly increases the perceived ease of use, perceived usefulness, and cognitive concentration"*, a conclusion complemented by the fact that tactile stimuli trigger generally shorter response times in users, compared to auditory and visual ones [6, 8]. Moreover, Hoggan et. al [6] conducted a user test where they compared the efficiency of tactile feedback of touchscreen keyboard with a non vibrating one, ultimately proving that haptic feedback enhances the experience by making the users faster at writing and more aware of typing errors even in loudly contexts such as the metro.

For the vast majority of the time, users interact with a small number of components, called widgets, while using the smartphone. A few examples are buttons, drop-down menus, toggles and sliders. Knowing that vibration is a fundamental aspect, a literature review related to the topic of smartphones' haptic feedback led to realising that none delved into it enough to provide guidelines for the tactile feedback design on widgets. With that in mind and since Android is the most popular mobile operating system in the world [1], we decided to focus on it to find out what are the most suitable patterns for each widget, whether there are differences between different Android platforms and, finally, provide design tips to the developers who want to approach this important part of the interaction. Therefore, the research question was defined as: "what are some rules of thumb developers should always follow, when implementing tactile feedback for Android widgets?".

For the purpose of answering it, we developed an app presenting several widgets with associated feedback and created a survey to evaluate them in a neutral context, which are presented in detail in the Method section of this report. After analysing the data collected and assessing what the trends were, as described in the Results section, some considerations and tips were elaborated and presented in the Discussion section, addressing the aforementioned target group.

BACKGROUND

Different information can be conveyed by modifying different attributes of a vibrotactile stimulus, as mentioned by Buzzi et al. [2]; therefore, alongside intensity of vibration, aspects such as gap length, number of gaps and vibration length can be focused on when developing a mobile application. These, in turn, influence the perception of urgency related to an event in the application, with an impact that decreases from the first to the last [9]. Outside the domain of mobile applications, signal frequency is another potential variable in producing haptic feedback: in a study by Van Erp and Spapé [5] it was found that up to nine different levels can be perceived by humans, and they evoke different feelings independently from the variation of the signal amplitude. As a comparison, the intensity of the vibration (corresponding to the signal amplitude) can only be perceived to a maximum of four different levels [9].

[1] https://gs.statcounter.com/os-market-share/mobile-tablet/worldwide/#monthly-201904-202004-bar, accessed on May 24, 2020.

Despite the human perception of tactile stimuli on their fingertips being extremely precise, thanks to the presence of a high number of mechanoreceptors [3], even in modern mobile devices vibrotactile feedback is not yet localised in a specific point of the device (or screen), which undermines accuracy in overall identification [2]. Rather, the entire display surface vibrates, and as an indirect consequence a significant noise is emitted: as Pomper et al. [8] suggest, the acoustical bias that might be caused can be overcome by using either loudspeakers or headphones in the context of the test. Moreover, more complex vibration patterns that could be achieved by, for instance, implementing high-fidelity force-feedback, are nowadays still not being included as a form of haptic feedback for most of the mobile device in the market [1].

More specifically, according to Choe and Schumacher intensity of vibration appears to have a positive impact on the second and third of these aspects. In Choe and Schumacher [3], one of the conclusions is that *"length of vibration is not a critical factor"*; still, with the aim of including long-lasting vibrations (e.g. for alerts or entertainment), the battery drain would be effected regardless, as the power draw of long-duration effects is significantly different from the short-duration ones [7]. On the contrary, Chu et al. [4] proposes that length of vibration *does* have an impact on the cognitive processing by the user when interpreting a long sequence of signals, but in a negative way: the more the pattern lasts, the more cognitively demanding it results.

All of the above mentioned considerations were used as a basis for the design of the Android application and its evaluation in the presented project. In particular, remarks from several different studies about the length of the vibration patterns in haptic feedback, as well as the minimum vibration length and the number of different vibrations that are perceivable by the average user were taken into consideration throughout the R&D process. Some of those were included amongst the design tips that are presented in the Discussion section for mobile developers.

METHOD

In order to assess the quality of a generic haptic feedback associated to a widget on a mobile application, three aspects were taken into consideration when interacting with it: how "natural" the vibrations feels; how "effectively" it is tied to the interaction itself; how much it catches the user's "attention". The first one indicates the measure of similarity between an action performed on a touchscreen and a corresponding one on a physical device. The second is based on the bond between the action and the expectations about its effects. The third stands for how easily noticeable the feedback is, i.e. in what measure it distracts the user from whatever else they are doing and focuses their attention on the widget.

Six widgets were selected among the native ones proposed on the Android Developers platform [2], as they are sufficiently different from each other and relevant from an interaction perspective (considering press, release, selection, deselection and slide as possible actions):

- *Button* → Pressing and releasing.
- *Seekbar* → Pressing, releasing and sliding.
- *Checkbox* → Selecting and deselecting.
- *Switch* → Selecting and deselecting.
- *Spinner* → Selecting.
- *Ratingbar* → Pressing and releasing.

All of them are illustrated in the first row of Table 1. Notice that they were not customised in any way, but rather used in their default implementation, so that the results are still applicable for every Android developer that decides to use them, regardless of how much they are modified afterwards. This means that only the listeners suggested on the aforementioned official developers platform were included in the code, and only the events related to the supported actions were handled.

Each widget comes in four different versions, with the first one being a control version and the others having different vibration patterns assigned, as shown in Table 1. Every pattern comprises of one or two vibrations, the length of which typically ranges from 4 to 100 milliseconds, but was sometimes set to 0 in order to exclude completely a particular action (or part of it). Pauses are encoded either with predefined value (50 ms) or depending on how long the user keeps contact with the widget.

With the aim of making sure that they made no accidental mistakes during the evaluation process, for both Android application and survey, participants were told that they were free to navigate them back and forth as well as interact with the widgets for as much as needed. At the same time, to avoid any kind of bias they were also advised to use headphones/earphones or listen to loud music during the test, so to cover any noise produced by the smartphone's vibrating motor.

Android application

The Android application itself was not without its own challenges and obstacles to face. First among those were the software limitations to vibration support in different platforms, that prevented the exploitation of some widgets' full potential: examples are the lack of handling for sliding actions by the default Ratingbar implementation, as well as for selecting options in the Spinner which are already selected, but also merely opening its drop-down menu. Additionally, a constant matter of attention was the appearance of the UI, which had to be as "neutral" as possible, implying that all versions of the same widget must have the same size, alignment, colour, as no added animations or sounds.

Development

The application was developed in Java on Android Studio 3.6.1 for Android 8.0 Oreo and tested on Samsung Galaxy S8 and S9+, starting from a template provided during a laboratory session of the Haptics, Tactile and Tangible Interaction course.

[2] https://developer.android.com/guide/topics/ui/look-and-feel/, accessed on May 24, 2020.

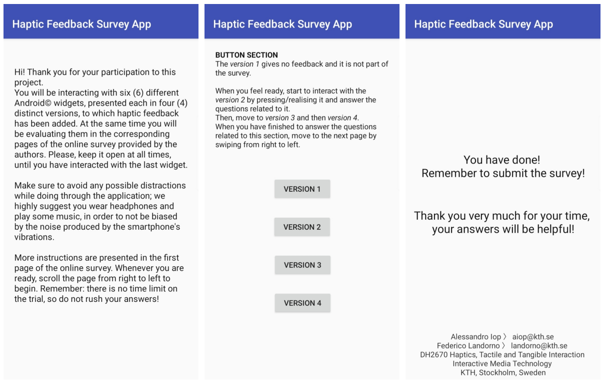

Figure 1. Screenshots of the three different kinds of fragment included in the application. From left to right: instruction page, Button page, final page.

The structure, as illustrated in Figure 1, is composed by eight different fragments that can be swiped left and right for navigating and include: a first page where the general instructions are displayed before beginning the test; six central fragments (one for each widget) showing the four versions in addition to an explanation of how to interact with them; a last page to thank the user for taking part in the test. Since the intention was to spread the Android application to an English-speaking as well as an Italian-speaking audience, the app is bilingual, meaning that the language used is programmed to be the one used in the user's smartphone.

Deployment

For deploying the application and spreading it to as many people as possible to test it, the overall process was kept fairly informal and direct – a far quicker solution than deploying it as it is meant to be. It was decided, in fact, to use its debug version, which is easier to set up as an `.apk` file, and then as many close friends and relatives as possible were contacted personally and given proper instructions on how to install it on their Android smartphones. It was enough to change the phone "third party applications" settings and enable installation of apps from the messaging service that the `.apk` file was sent through, then by just clicking on the file received, the app was downloaded and installed. Calling almost everybody took a relevant amount of time, but delivering clear explanations was ultimately worth it. In particular, when entire family units were involved in the evaluation of the application, having only one smartphone to test from was recommended, as it was easier to manage and accessible to non-Android users as well.

In summary, then, reaching a wider pool of users for the evaluation test, including (possibly) also iOS users, was deemed to be more important than making sure that the testing environment was totally controlled and organic across all participants. Collecting feedback from people accustomed to different platforms could, as a matter of fact, potentially provide interesting insights into the subject matter of the current project. Despite the deployment of the application being somewhat informal under a few aspects, consistency was kept across all evaluations by never changing the content of the app in any circumstances, even if some better instructions could have been written.

Survey

The main reason behind the choice of spreading an online survey as tool for conducting the evaluation were the social distancing recommendations given during the COVID-19 pandemic, which prevented us from setting it up in a more controlled environment. In addition to that, having potentially a bigger reach on the audience supported this choice, as it is generally easier to get a higher number of participants this way. The survey itself was created using Microsoft Forms in two languages (Italian and English), as the application itself was (see the Development subsection), so to reach a wider pool of users. It was meant to be filled in while interacting with the widgets on the app, and therefore participants were given the freedom of keeping it open either on the same smartphone the app was installed on, on another smartphone or on a computer.

Instructions on how to fill the questionnaire in and how to install the application were presented in the first two pages, which were followed by several general questions including: age group; gender; whether the participant was using their own smartphone in addition to, in case they did not, what operating system their own smartphone has; the brand of the smartphone they were currently using for the evaluation and, if known, its model name. Such questions were included with the aim of taking them into consideration when addressing potential outliers in the data that would be collected from the survey responses, and with the assumption that all of it would be successively treated completely anonymously. For this reason, each one of the core pages contains a privacy disclaimer warning testers:

> *Please note that the present survey is entirely anonymous and the data will be collected for internal use only. The information you will be submitting will not be shared with 3rd parties, in compliance with the GDPR regulations.*

The remaining 6 pages of the questionnaire were associated to the different widgets under analysis; in particular, for each of them a set of three questions related to the aspects presented in Method section was proposed, as the following ones, concerning the Button, exemplify:

1. How "natural" did the vibration feel when interacting with the Button, in a scale from 1 to 5?
2. How effectively is the vibration tied to the specific action of interacting with the Button, in a scale from 1 to 5?
3. How much does the vibration catch your attention when interacting with the Button, in a scale from 1 to 5?

For every one of them, a grid of Likert-style ratings, corresponding to each version of every widget, allowed values on a scale from 1 (one) to 5 (five) as well as a "not applicable" option. The latter was included because it was expected that the first version of any single widget (which was the one *without* any tactile feedback), would be marked as such under any of the three aspects considered. As mentioned in the Discussion section, it would not be so. A snippet of one of the core pages of the survey is displayed in Figure 2.

Addressing that the meaning of a question can be interpreted differently by each participant and depends on the choice of

Widget name	Button	Seekbar	Checkbox	Switch	Spinner	Ratingbar
Image				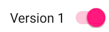	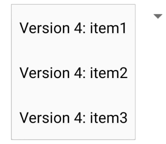	
Version 2 (ms)	10 – 0	Progress×1.5	10 – 0	30 – 0	Item index×20	Rating×10
Version 3 (ms)	10 – 10	4	30 – 10	10 – 10	50	100
Version 4 (ms)	30 – 10	4 when pressed and released	10 – 30	30 – 10	50 – 50 – 50 – 50	Rating×40

Table 1. Tactile feedback patterns applied to every version of each widget, in milliseconds. *Progress* and *rating* depend on the current values of the Seekbar and Ratingbar, respectively. Version 1 has never any haptic feedback assigned, therefore it is omitted from this table.

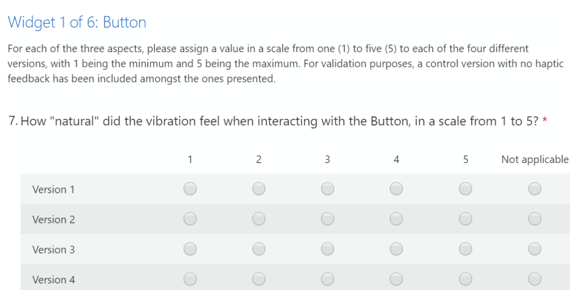

Widget 1 of 6: Button

For each of the three aspects, please assign a value in a scale from one (1) to five (5) to each of the four different versions, with 1 being the minimum and 5 being the maximum. For validation purposes, a control version with no haptic feedback has been included amongst the ones presented.

7. How "natural" did the vibration feel when interacting with the Button, in a scale from 1 to 5? *

	1	2	3	4	5	Not applicable
Version 1	○	○	○	○	○	○
Version 2	○	○	○	○	○	○
Version 3	○	○	○	○	○	○
Version 4	○	○	○	○	○	○

Figure 2. Screenshot of the upper part of a specific widget’s page in the survey, presenting instructions for filling in the forms at the top and the first of the three questions at the bottom.

words to express it, an open-ended, non-mandatory question for submitting feedback about the specific associated widget was included at the end of the page. This would hopefully ensure that all other considerations with regards specific details about the test were collected. To accomplish this, users were also suggested to fill in the page corresponding to a specific widget right after interacting with its different versions in the application.

RESULTS

The evaluation process was conducted among 22 participants, of whom 14 were males and 8 females, ranging from 18 to 60 years of age; with the exception of one tester, all of them own an Android smartphone – which means that they are used to interacting with the proposed interface. 14 participants tested the application with their own device while the other 6 used a borrowed one. All of their responses were collected on a Google Sheets[3] spreadsheet, which can be found in the Resources section and where the six widgets are divided among as many sheets, containing raw data along with graphs generated for statistical analysis. On a separate sheet, the main results and observations pertaining all aspects of every widget are presented as a table that takes different graphs into consideration separately. Three kinds of charts were included:

- a stacked column chart (the Button, for example, is depicted in Figure 3) with *absolute frequencies* for each possible value, including “Not applicable”, which was chosen because it visually summarises all data in one place and enables some early observations to be drawn;
- a column chart (the same Button example is in Figure 4) grouping widget versions for each one of the three aspects considered – i.e. “Natural-ness”, effectiveness and attention grabbed – without taking “Not applicable” responses into account, which allows comparisons between the *average* values of different tactile feedback patterns that are placed side by side;
- a candle chart (Figure 5, again for the Button example), displaying the *distributions* of all non-“Not applicable” responses, sorted by the aspect they are associated to, in a way that is similar to a boxplot.

The first version of each widget, which has no feedback associated to it, was not considered during the analysis of the results because it was intended as a control version. Therefore, for the rest of this section every time “all versions” is mentioned, only version 2, 3 and 4 are considered. The only relevant cases involving the version 1 are the amount of attention it grabbed for the Button, according to the column chart, as well as how every aspect scored for it, according the Ratingbar’s candle chart.

A peculiarity is that across almost all the responses the users gave, the whole range of values in the Likert-style rating scale is covered. More specifically, almost every aspect of every widget version is rated from a minimum of 1 (one) to a maximum of 5 (five).

Button

All versions are almost equally natural and effective based on column and candle charts. Concerning the attention aspect, responses for version 2 exhibit a heavy distribution around

[3] https://www.google.com/sheets/about/, accessed on May 28, 2020.

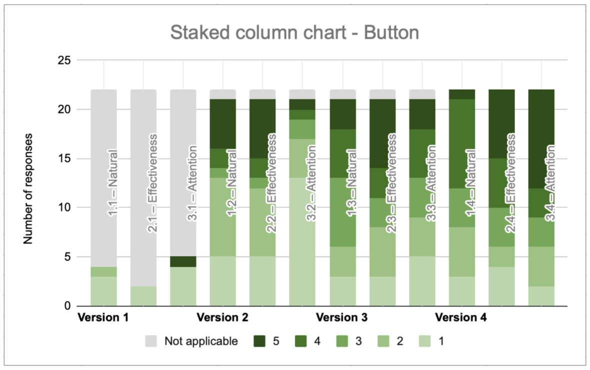

Figure 3. Stacked column chart for the Button widget. Distribution of the 22 total responses across all possible ratings (including "Not applicable", in grey), mapped to the vertical axis, with shades of green corresponding to a higher value. The columns are grouped by version and sorted by aspect (first "natural-ness", then effectiveness and finally attention).

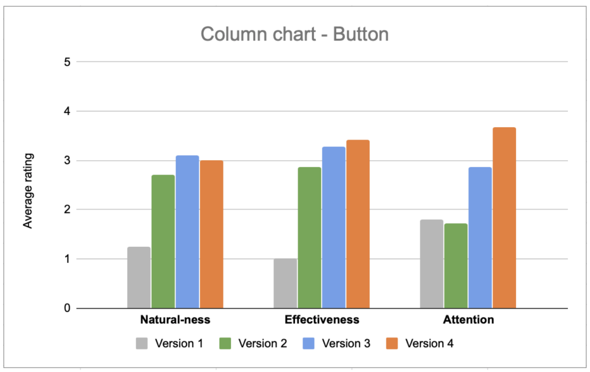

Figure 4. Column chart for the Button widget. Average of the rating given by the participants (excluding "Not applicable"), grouped by aspect and sorted by version.

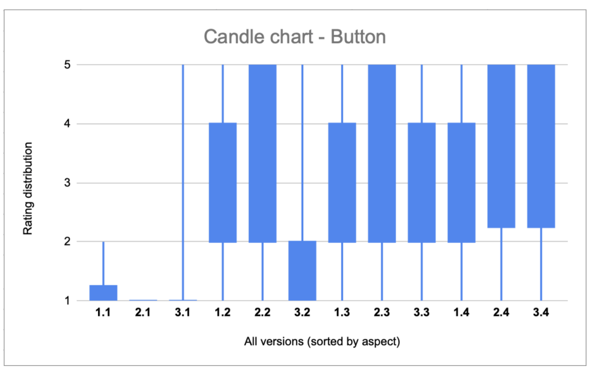

Figure 5. Candle chart for the Button widget. Variation of the rating values for all aspects of each version (excluding "Not applicable"). Candles are grouped by version and sorted by aspect, like in the stacked column chart, meaning that for every label on the horizontal axis the first number corresponds to a specific aspect (1: "natural-ness", 2: effectiveness and 3: attention grabbed) and the second number corresponds to the version.

their medians. Additionally, according to all three graphs the amount of attention grabbed is considerably lower for the second version (even lower than for the first one) while the stacked column chart, in particular, indicates that the fourth grabbed the most. Thus, there seems to be a correlation between an higher amount of attention grabbed and the complexity of the vibration pattern, where a pattern with more or longer vibrations is defined as more complex than one with fewer or shorter ones. Such short-lasting vibrations as the ones associated to versions 2 and 3 often happened to be barely (if at all) perceived by a good portion of the participants – six at least – who left a feedback in the Button's page of the survey. This suggests either that single vibrations lasting around 10 milliseconds or less are not supported by some of the smartphones used during the evaluation, or simply that the resolution of the mechanoreceptors on the users' hands is too low for the stimuli to be felt.

Seekbar

While according to the column chart and the stacked one, as noticed for the Button, higher pattern complexity corresponds to more attention being caught by the haptic feedback, it also corresponds to a lower average "natural-ness" across all the versions. The effectiveness of this widget, on the other hand, does not appear to vary in a relevant way. Another observation in common with the Button is the inability of a few of the participants to perceive one or more different feedback or barely so, as reported in their comments. Most frequently the fourth version was judged to be lacking any stimuli, due probably to the short length (4ms) of the vibrations associated to it, which are only produced when the slider is touched for the first time and when it is released.

Checkbox

With regards to the first and third of the aspects under analysis, according to the column chart as well as the stacked one, the same observations can be made on the relation between pattern complexity and their ratings. Upon looking at the candle chart, instead, it can be noticed that for "natural-ness" versions 2 and 4 are equally distributed, while version 3 is more loosely distributed; at the same time, for attention all versions have different distributions, with version 4 having the loosest one. From the very same graph, another item of interest is that versions 2 and 4 are equally distributed under an effectiveness standpoint, while version 3 is more concentrated. Similarly to the previous widgets presented, a set of comments expressed inability to perceive or distinguish one or more patterns.

Switch

It can be argued that all versions have almost the same average rating and similar distributions around their medians, especially for "natural-ness" and effectiveness. With a lesser degree of confidence, it can also be hypothesised that version 4, which has the slightly more complex vibration pattern, grabbed more attention but, since the other versions have loose distributions, it is difficult to support. Feedback from a few of the participants seems to confirm the former argument in particular: sometimes vibrations were not perceived when the Switch was turned off, and/or all versions were felt the same.

Spinner

Versions 2 and 3 of the Spinner are similar considering all the aspects, both in the stacked column chart and in the column chart, however the same cannot be assessed for the candle chart. In particular the second one has the loosest distribution around its median when considering the amount of attention it gets, which might be related to the wide range in length of the vibrations in its pattern. The participants found version 4 (the one with a pattern of two vibrations of 50 milliseconds separated by the same amount of time) out of context because the given feedback provided more stimuli than the inputs performed on the touchscreen. This version had also the lowest averages across both “natural-ness” and effectiveness.

Ratingbar

For the Ratingbar, “natural-ness” and effectiveness are almost equal for version 3 and 4 while the second was perceived to be remarkably more natural and there seems to be a correlation between the range of vibration length and the the distribution of its ratings: the more the vibration vary in length, the more loosely distributed the versions are. Concerning attention, it can be noticed that higher pattern complexity corresponds to more attention grabbed on average.

DISCUSSION

From the results gathered and analysed, some relevant trends can be extrapolated for the different widgets and they can be summarised in the following five trends.

1. Very short vibrations, especially when associated to a single brief action (e.g. press and release), are often difficult to perceive by the users, because of their biophysical limitations, or impossible to be performed by some Android devices. This is highlighted in the Button, Seekbar (version 4 in particular), Checkbox and Switch subsections.
2. As observed in the Switch subsection, similar haptic feedback patterns are associated to a similar distributions and average ratings when considering “natural-ness” and effectiveness. By performing more statistical analysis on the data, this could also be extended to the other widgets.
3. Lower levels of “natural-ness” are experienced for higher complexity patterns (see Seekbar and Checkbox subsections), where the complexity itself is defined by the number of vibrations, number of gaps, and/or their respective lengths.
4. Pattern complexity, in turn, is directly proportional to the amount of attention the tactile feedback gets from users; an argument supported in the Button, Seekbar, Checkbox, Spinner (in particular version 4) and Ratingbar subsections.
5. Transitively, from points 3 and 4 it can be entailed that there is an inverse relation of proportionality between “natural-ness” and attention, meaning that lower perceived levels of the former correspond to higher ones of the latter.

According to Banter [1], an output vibration that depends directly on the value of a parameter of the object – in this case, a widget – it is associated to ultimately enhances the user experience. This approach was tested on Seekbar (version 2), Spinner (version 2) and Ratingbar (versions 2 and 4) by setting the haptic feedback in relation to progress, number of the items and rating respectively, but only the last widget provided a result that agrees with Banter’s statement. Indeed, it was found that the participants did not notice any differences between a static pattern and an incremental one. Similarly, a strong expectation we had was that the responses about the first version of each widget would be deemed “Not applicable”, since no tactile feedback was applied on them. Instead, not only some of them got low values – which is understandable, after all – but a few even got high ones, thus contradicting the expectations, such as for the Button’s attention, the Checkbox’s effectiveness, the Switch’s “natural-ness” and effectiveness as well as all Ratingbar’s aspects. Naturally, some of the higher ratings could have been given by participants by mistake, however it is impossible to say with certainty; concerning those cases where 3 (three) was given, it might mean that a specific version was perceived as neutral for specific aspects. On the other hand, the amount of occurrences where all the other versions, for each widget, were deemed “Not applicable” is always negligible and therefore excluding this option from the column and candle charts does not have any significant impact on the overall results. This is why the conclusions drawn above can be considered relevant in the context of the evaluation study presented here.

Another observation that can be made concerns the distribution of every version of every widget under any of the aspects that are being focused on, which is rarely very concentrated around the median and, as mentioned in the Results section, almost always covers all possible values, provided that version 1 is not considered. That is, the data points are very often spread on a wide range of possible values, as shown in Figure 5, for instance. This might be due to the small number of participants that responded to the online survey, therefore to be able to have shorter “candles” in the graphs a possible improvement of the present study could be reaching more people for the evaluation. This, in turn, would constitute a more solid basis for the observations made above. Alternatively, a more controlled testing environment could be set up instead of spreading the survey among friends and family members: more precise instructions, less bias due to environment distractions, same Android smartphone(s) used by all participants, a lighter cognitive load by not having to download and install the application before the test are but a few of the possible design choices that can be made to polish the evaluation.

In addition to shorter candles, the basis for the conclusions made so far could be further strengthened by performing a more in-depth statistical analysis that would provide useful confidence intervals. Furthermore, introducing more versions for fewer widgets would also focus the study on a wider variety of patterns for a smaller pool of components that potentially combines a higher number of parameters. These could be, for instance, lengths of vibrations and gaps between them, their variance, their range of length and the correlation with the widget’s status or value. Finally, conducting a proper pilot test – which only involved the two developers for the present project – with external people would prevent any potential fallacies caused by the Android widget handlers and events that are not supported by certain platforms and devices.

Upon bringing all of the aforementioned discussion points together, several tips can be come up with to aid developers who might want to introduce tactile feedback to Android applications. They are meant to be general rules of thumb that should be followed when programming the interaction with the six Android widgets considered in this study; according to the specific context, or domain of application, they may differ slightly from their original version. The 8 tips are the following:

- use more than 10ms for single vibrations;
- for sliding actions you can use vibrations even shorter than 10ms;
- seek correspondence between the number of actions performed and of vibrations in the pattern;
- avoid vibration lengths that are proportional to the widget's value, because they have no relevant impact on the user experience;
- use reasonably short vibrations at all times;
- don't give different lengths to vibrations in the same pattern, because they will pass unnoticed;
- always verify that the vibrations are long enough to be perceived, by testing the UI;
- check that all interaction events are managed as expected, and their listeners are supported by all platforms.

In conclusion, this research presented what are the features of the haptic feedback that make it more suitable for each of the six widgets considered, assess that there exist differences in what measure different Android devices support tactile patterns which should be accounted for even when designing them, and provides rules of thumb for developers to follow in order to ensure that the user experience is improved by the haptic feedback.

RESOURCES

The full spreadsheet, divided into tabs (one for each widget evaluated) and complete with the relative charts, can be accessed at `this link`. The code of the Android application can be found in `this GitHub repository`. The slides of the project presentation, held on May 19, are available `here`.

REFERENCES

[1] Bruce Banter. 2010. Touch Screens and Touch Surfaces are Enriched by Haptic Force-Feedback. *Information Display* 26, 3 (2010), 26–30.

[2] Maria Claudia Buzzi, Marina Buzzi, Barbara Leporini, and Caterina Senette. 2015. Playing with geometry: a multimodal android app for blind children. In *Proceedings of the 11th Biannual Conference on Italian SIGCHI Chapter*. 134–137.

[3] Pilsung Choe and Dennis Schumacher. 2015. Influence of different types of vibrations on technical acceptance of a mobile game aiming for hedonic satisfaction. *International Journal of Human-Computer Interaction* 31, 1 (2015), 33–43.

[4] Shaowei Chu, Fan Zhang, Naye Ji, Fangyuan Zhang, and Ruifang Pan. 2017. Experimental Evaluation of Tactile Patterns over Frictional Surface on Mobile Phones. In *Proceedings of the Fifth International Symposium of Chinese CHI (Chinese CHI 2017)*. Association for Computing Machinery, New York, NY, USA, 47–52. DOI:http://dx.doi.org/10.1145/3080631.3080639

[5] Jan B. F. Van Erp and Michiel M. Spapé. 2003. Distilling the underlying dimensions of tactile melodies. Proceedings of Eurohaptics. (2003), 111–120.

[6] Eve Hoggan, Stephen A. Brewster, and Jody Johnston. 2008. Investigating the Effectiveness of Tactile Feedback for Mobile Touchscreens. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '08)*. Association for Computing Machinery, New York, NY, USA, 1573–1582. DOI: http://dx.doi.org/10.1145/1357054.1357300

[7] Michael Levin and Alfred Woo. 2009. Tactile-Feedback Solutions for an Enhanced User Experience. *Information Display* 25, 10 (2009), 18–21.

[8] Ulrich Pomper, Jana Brincker, James Harwood, Ivan Prikhodko, and Daniel Senkowski. 2014. Taking a call is facilitated by the multisensory processing of smartphone vibrations, sounds, and flashes. *PloS one* 9, 8 (2014).

[9] Bahador Saket, Chrisnawan Prasojo, Yongfeng Huang, and Shengdong Zhao. 2013. Designing an Effective Vibration-Based Notification Interface for Mobile Phones. In *Proceedings of the 2013 Conference on Computer Supported Cooperative Work (CSCW '13)*. Association for Computing Machinery, New York, NY, USA, 149–1504. DOI: http://dx.doi.org/10.1145/2441776.2441946