Smartphone Usage Adaptation to Phone Battery level

Bram de Kroon
STUDENT NUMBER: 2027371

THESIS SUBMITTED IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS FOR THE DEGREE OF
MASTER OF SCIENCE IN COGNITIVE SCIENCE & ARTIFICIAL
INTELLIGENCE OR DATA SCIENCE & SOCIETY
DEPARTMENT OF COGNITIVE SCIENCE & ARTIFICIAL INTELLIGENCE
SCHOOL OF HUMANITIES AND DIGITAL SCIENCES
TILBURG UNIVERSITY

Thesis committee:

Drew Hendrickson
Merel Jung

Tilburg University
School of Humanities and Digital Sciences Department of Cognitive
Science & Artificial Intelligence Tilburg, The Netherlands
December 2019

Preface

This thesis was written in partial fulfillment for the requirements of the degree of Data Science & Society. I would like to thank Drew Hendrickson and Giovanni Cassani for their valuable guidance while conducting the research.

Smartphone Usage Adaptation to Phone Battery Level

Bram de Kroon

Mobile phone technologies have developed rapidly over the past few years. To be able to facilitate these continuously developing technologies, smartphones demand increasingly more battery capacity. Improving smartphone energy efficiency is an ongoing challenge and is being addressed from numerous perspectives. On another note, there is increasing interest in understanding how smartphone users use their phone. This work is an initial attempt to determine how smartphone users adapt their phone usage behavior to the battery level of their phone. Answers to this matter might prove to be relevant for research regarding smartphone energy efficiency as well as the understanding of smartphone usage itself. Six features have been evaluated which all represented a concept of phone usage behavior and quantified to what extent a smartphone user adapts the respective concept of phone usage behavior given two intervals of battery level. Results suggest that distinct patterns of change in phone usage behavior cannot be accurately captured using global intervals of battery level. Instead, they suggest that we should look closely to how smartphone users adapt their phone usage behavior to a more continuous scale of battery level.

1. Introduction

Mobile phone technologies have developed rapidly over the past few years. State-of-the-art hardware boosts the computational power of mobile phones and enables the development of more sophisticated mobile applications. Simultaneously, wireless technologies continue to develop and the imminent introduction of 5G networks promises to push the capabilities of smartphones even further. To be able to facilitate these continuously developing technologies, smartphones demand increasingly more battery capacity. To date, extending the battery life of smartphones is an ongoing challenge that is being addressed from various fields of research.

Battery life is generally referred to as run time on a full battery charge. Research efforts which aim to extend battery life can be subdivided into two broad categories, namely the improvement of energy efficiency and the improvement of battery capacity. The latter concerns the development of the battery itself, for instance, the development of next generation Li-ion batteries (Deng, 2015). This category shall not be further discussed as it falls outside of the scope of this research.

The ongoing challenge of improving smartphone energy efficiency is being addressed form numerous perspectives. Recent work focused on understanding the interactions between batteries and applications (Guo, Wang & Che, 2017), battery level predicting (Li, Liu & Mei, 2018; Oliver & Keshav, 2011), energy modeling schemes (Yurur, Liu & Moreno, 2015; Ahmad et al., 2017), smartphone charging habits (Ferreira, Dey & Kostakos, 2011) and more. However, to my knowledge, little to no efforts have been made to understand how people adapt their smartphone usage behavior to the battery level of their phone.

Knowledge of smartphone usage adaptation is interesting from a scientific point of view as it might serve as an additional dimension in research regarding the understanding of phone usage behavior in general. Moreover, this knowledge can potentially be exploited by research fields concerned with the development of smart energy management systems (Datta, Bonnet and Nikaein, 2014; Draa, Niar, Tayeb, Grislin & Desertot, 2017) as it might encompass relevant information which can lead to battery savings. Development of such systems is interesting for mobile phone developers as battery life is an important feature that directly affects the popularity of smartphones (Kekolahti, Kikki, Hämmäinen & Riikonen, 2016). Additionally, increasing smartphone battery life could stimulate the increase in availability of smartphone log data as battery drain is one of the main concerns that withholds people from participating in such data collection efforts (Anjomshoa & Kantarci, 2018).

This project is an initial attempt to determine how smartphone users adapt their phone usage behavior to the battery level of their phone. The aim of this project is to ascertain whether smartphone users adapt their phone usage behavior to phone battery level in different ways. In addition, the project aims to assess whether smartphone users adapt their phone usage behavior gradually or more erratic across levels of phone battery. By extracting features from smartphone log data, this thesis intends to answer the following research questions:

- To what extent can we distinct different patterns of phone usage adaptation to phone battery level by clustering smartphone users based on features that describe how smartphone users adapt their phone usage behavior to intervals of battery level?
- Are the defined features and cluster solutions stable when we evaluate how smartphone users adapt their phone usage behavior considering different intervals of battery level?

To be able to answer these questions, features have been engineered that describe how smartphone users adapt general notions of phone usage behavior given two intervals of battery level. Initially, feature values have been computed considering battery levels 1-50 and 51-100 as intervals. The features quantify to what extent a participant exhibits less or more of the respective concepts of phone usage behavior at the higher interval of battery level (e.g., if the user uses applications for shorter or longer periods of time when phone battery level is 51 percent or higher, as compared to when phone battery level is 50 percent or lower). The features have been clustered using the DBSCAN algorithm. The cluster results have been evaluated by means of plots and the density-based clustering validation metric (DBCV).

The top 10 cluster results in terms of the DBCV measure all consist of one cluster. With a minimum DBCV measure of 0.709, all evaluated cluster solutions are of considerable density. Results indicate that smartphone users do not exhibit clear distinct patterns of phone usage adaptation to battery level. Instead, they suggest that phone usage adaptation fluctuates more gradually between users. However, further analyses of the used features unveiled large differences in feature values when considering more narrowly defined intervals of battery level. In addition, the cluster solutions were found to be rather unstable when comparing them to cluster solutions that were established using features that have been computed over the more narrowly defined intervals of battery level. These results indicate that the approach using battery levels 1-50 and 51-100 as intervals is too broad, and that phone usage adaptation might be best captured by analyzing change in phone usage behavior given multiple, more narrowly defined intervals of battery level.

2. Related Work

Understanding how smartphone users adapt their phone usage behavior to the battery level of their phone has, to my knowledge, not been addressed by the literature. However, it does overlap with and can be found relevant in two areas of research, namely the field of research concerning energy efficiency and the field of research concerning phone usage behavior in general. Both fields are closely intertwined given that phone usage behavior is often used as input in various works regarding battery efficiency. Important facets from both fields will be discussed in this section.

The increasing availability of smartphone log data facilitated researchers to gain a deeper and more accurate understanding of how smartphones are being used. Smartphone user diversity has gained a lot of attention and illustrates how smartphones are being used across various types of users.

One of these works analyzed how smartphone usage statistics differ across frequent and less frequent smartphone users (Guo, Wang & Chen, 2017). The authors classified participants as heavy, normal or lights smartphone users based on phone usage duration time. In this work, the 20% most frequent smartphone users were classified as heavy users and the 20% least frequent smartphone users represented the light class. Analysis unveiled vast differences in daily smartphone usage duration times between the heavy group (8.6 hours per day) and the light group (0.9 hours per day).

In a more comprehensive analysis, Falaki et al. (2010) found similar and additional differences in user diversity. Their work showed that average daily smartphone usage time varies from 0.5 to 8.3 hours. In addition, the amount of daily smartphone usage sessions (varying from 10 to 200) and session length (varying from 10 to 250 seconds) differ greatly between users and were found to be uncorrelated.

Complementary to the more general notions of smartphone usage, studies also attempted to understand how users interact with the vast amount of available smartphone applications. Guo et al. (2017) found that applications used for communication and internet browsing accounted for approximately 70% of smartphone usage time. This finding is relatively consistent across other works (Brown, McGregor & McMillan, 2014; Falaki et al., 2010). Moreover, Guo et al. (2017) found application usage to be extremely diverse and yet very concise: their data set contained roughly 23,000 different applications of which the 200 most frequently used applications accounted for 82% of total smartphone usage time. Neither Guo et al. (2017) nor Falaki et al. (2010) found significant differences between the distributions of used application categories of heavy and light smartphone users.

More variation arises as we look at usage times of application categories. Falaki et al. (2010) found that application categories differ from each other in terms of mean usage times. In addition, they found that smartphone users distinct themselves from others by using similar applications for different amounts of time. Ferreira, Goncalves, Kostakos, Barkhuus and Dey (2014) studied application micro-usage, which they defined as “brief bursts of interaction with applications” (p. 91). Their work provides insight into the proportions in which applications are used for short and longer periods of time. Findings indicate that these proportions differ across application categories.

It is evident that smartphone usage behavior is extremely diverse. Various works have established differences in smartphone usage but fall short in explaining this diversity. Nevertheless, an understanding of smartphone usage behavior is of importance for research regarding smartphone battery usage as phone interaction (Datta et al., 2014) and application usage (Guo et al., 2017) considerably affect battery lifetime.

In addition to diversity in smartphone usage, diversity has also been found in how users interact with phone batteries. In an early work, Banerjee, Rahmati, Corner, Rollins and Zhong (2007) studied battery use and battery recharge behavior of laptop and mobile phone users. The authors argued that accurate knowledge of this matter can be exploited to develop more sophisticated energy management systems. Ferreira et al. (2011) conducted a similar kind of study focusing solely on charging behavior of smartphone users. The authors found that users tend to keep battery levels above the 30%

mark. Moreover, their results unveiled two types of charging behavior: arbitrarily charging batteries for short periods of time and charging batteries to maximum capacity. Oliver and Keshav (2011) identified users who charge smartphone batteries for multiple short periods of time to be the most excessive battery consumers.

So far, we discussed how users interact with their smartphone and its battery. This information is closely linked to the task of predicting battery lifetime. In general, research concerning battery lifetime prediction aims to enable smartphone users to effectively plan phone activities by providing them with information regarding how long the battery will last. The task of predicting battery lifetime is inherently complex as it is affected by numerous variables such as hardware, software and phone usage itself (Li et al., 2018).

Works attempted to predict battery lifetime from different perspectives. A relatively consistent factor in these works is the use of hardware status as predictors. Works distinguish themselves by using additional predictors such as smartphone usage patterns (Chantrapornhai & Nusawat, 2016; Kang, Seo & Hong, 2011), application usage (Kim, Chon, Jung, Kim & Cha, 2016) or a more comprehensive approach (Li et al., 2018).

In a simplified approach, Kim et al. (2016) proposed a framework which predicts battery lifetime assuming that smartphones constantly run the same application. Using the status of various hardware components, the framework is able to predict battery lifetime with roughly 93% accuracy. Trying to predict battery lifetime based on more natural phone usage behavior, Li et al. (2018) extracted and evaluated the predictive performance of features from the Sherlock data set (Mirsky, Shabtai, Rokach & Shapira, 2016). Ultimately, their model was able to predict battery lifetime with an average error of approximately 107 minutes. Li et al. (2018) found historic battery consumption rate to be the most useful predictors given the prediction problem.

A variety of works have developed models that attempt to predict battery lifetime. Under simplifying assumptions, battery lifetime can be predicted with considerable accuracy, such as the work of Kim et al. (2016). However, accuracy tends to drop as works attempt to predict battery lifetime given more natural smartphone usage behavior (Li et al., 2018). It will be interesting to see if and how understanding of smartphone usage adaption to phone battery level will be able to enhance performance in such prediction tasks.

An exciting product that relates to the discussed areas of research is the development of energy management systems that adapt their policies to the specifics of smartphone users. As early as in 2007, researchers emphasized the potential of these kinds of systems (Banerjee et al., 2007; Ferreira et al., 2011). Datta et al. (2014) proposed an architecture that constructs phone usage patterns by monitoring smartphone usage. Using the constructed patterns, it then generates and implements a user-specific power saving profile. Evaluations showed that the architecture results in larger battery saving compared to well-established applications such as Juice Defender. More recent work attempted to implement real-time power saving profiles by predicting which applications a user is likely to use next (Draa et al., 2017) and tried to incorporate battery lifetime demands of users (Draa, Niar, Grislin-Le Strugeon, Biglari-Abhari & Tayeb, 2019).

An apparent conclusion is that smartphone usage is extremely diverse. The literature has established the diversity but falls short in explaining how it arises. The diverse behavior facilitates numerous types of analysis such as predicting battery lifetime. However, due to this diversity, we must be careful interpreting results as findings based on small sample sizes are likely to be prone to some kind of bias (Zhao et al., 2016) and are difficult to generalize (Church, Ferreira, Banovic & Lyons, 2015).

This work is an initial attempt to understand if and how smartphone users adapt their phone usage behavior to the battery level of their phone. One of the steps in this work is to find behavioral features which can be used to represent the concept of change in phone usage behavior. Doing so, it might provide valuable information which can be used to explain a portion of the large amount of diversity in smartphone usage.

Additionally, this work tries to assess whether smartphone users show distinct patterns of change in phone usage behavior. Features capturing these patterns might prove to be very useful in tasks such as predicting battery lifetime as they encompass information regarding how the smartphone user is going to adapt his phone usage. For instance, features describing how a user is going to adapt phone usage behavior as battery level decreases might significantly interact with existing features and, as a result, substantially increase real-time battery lifetime prediction accuracy. In a similar sense, the same features might also provide more ground for energy management systems to interact with and anticipate on user behavior.

3. Experimental Setup

This section contains a detailed description of the used data set, the used features and the adopted methodology and used models.

3.1 Data

The analysis was conducted on an existing data set containing mobile phone log data (Hendrickson, Aalbers & Vanden Abeele, under review). The data has been collected from 124 participants over the course of four weeks. The data has been logged through an application which participants installed on their personal smartphone. Observations in the data set represent the use of a smartphone application and contain information regarding the context of its use. Every observation encompasses information regarding which application has been used, the start and end time of the usage event, the battery level at the start of the event, a session number and a user-id. Every participant has a unique user-id. Smartphone usage sessions are considered to happen when multiple applications have been used consecutively. Therefore, observations representing consecutively used applications are given the same session number.

The unprocessed data set consists of 586,792 observations collected from 124 participants. Data exploration unveiled some problems in the data which have been solved as follow:

- In some cases, the same events have been recorded twice. All duplicate observations have been discarded.
- All observations with a recorded battery level of zero have been discarded. Observations with zero battery level were also likely to be outliers in terms of usage duration.
- Participants (4) have been removed as a result of corrupted session registrations. The respective participants were outliers in terms of session duration which was caused by an unexplainable large number of observations having similar session numbers.
- Participants (2) with no registered observations below 40% battery level have been discarded.
- Participants (11) who participated less than four days in the data collection effort have been discarded in an attempt to bolster the representativeness of the registered phone usage behavior.
- Application usage events that lasted shorter than half a second or longer than one hour have not been considered in this study. A considerable amount of application usage events (32,194) lasted shorter than half a second. However,

with a mean duration of 0.15 seconds, these events have not been considered meaningful. A small number of observations lasted longer than one hour (283). With a mean duration value of 5,860.69 seconds, these observations have been considered as outliers.

The cleaned data set consists of 513,799 observations collected from 104 participants. On average, participants participated 24.4 days in the data collection effort.

The data set has been further explored to get insight into how battery level affects smartphone usage. Approximately 62.2% of all application usage events have taken place while the battery level was 50% or higher. This difference can be partially explained by the literature as Ferreira et al. (2011) found that users have the tendency to keep battery levels above the 30% mark. However, it might also imply that users simply interact less frequent with their phone at lower battery levels.

Looking at how application usage events are distributed across battery levels, we find that a significant number of instances take place at 100% battery level (figure 1). This finding indicates that users keep their phones connected to a power source, even when the battery is fully charged. This is also supported by the work of Ferreira et al. (2011), who found that users overcharge phone batteries on a regular basis.

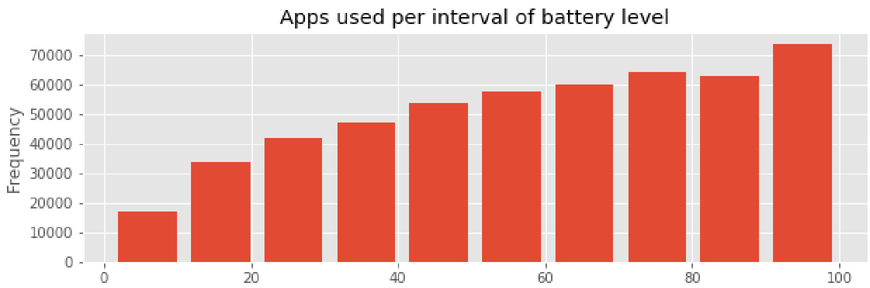

Figure 1: Bar plot visualizing the amount of applications used at different intervals of battery level

The mean application usage time across battery levels is visualized in figure 2. What is interesting is that, with a mean of 64.83 (*SD* = 192.82), users use applications for longer durations of time at battery levels below the 50% mark as compared to battery levels above the 50% mark (*M* = 60.13, *SD* = 180.14). A plausible explanation could be that smartphone users refrain from using applications for less urgent causes at lower battery levels.

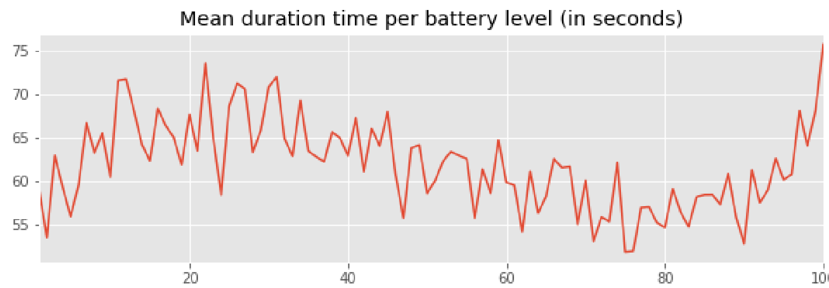

Figure 2: Line plot visualizing the mean application duration time at different values of battery level

It comes to little surprise that the data set contains highly variable information. Table 1 presents daily phone usage statistics that have previously been described in the literature (Falaki et al., 2010; Guo et al., 2017). Using the same heuristic as Guo et al. (2017), participants have been classified as heavy, medium or light phone users based on their daily phone usage time. The statistics show large differences between user types. On average, heavy users spend more than six times as much time interacting with their phones as compared to light users. This difference mainly results form the number of apps used (331.62 vs 71.01) as the difference in app usage duration (80.52 vs 53.59 seconds) is relatively small. It appears that medium and light users use an approximately similar number of apps per phone usage session while heavy users tend to use more apps during sessions.

	Phone usage (hours)	Apps used	App usage (seconds)	Sessions	Session duration (seconds)
Heavy	6.45	331.62	80.52	74.74	419.95
Medium	3.35	199.39	67.09	64.73	232.49
Light	0.94	71.01	53.59	24.19	166.34

Table 1: Descriptive statistics of daily phone usage. The top and bottom 20% of participants in terms of daily phone usage time have been classified as heavy and light users respectively.

When grouping the data based on battery level intervals, we see that heavy users show more consistent smartphone usage behavior across battery level intervals (table 2). Meanwhile, medium and light users use tend to use their phones less at the lower interval of battery level. All user types use their phone in longer sessions and use apps for longer durations of time at the lower battery level interval.

	Battery interval	Phone usage (proportion)	Apps used (proportion)	App usage (seconds)	Sessions (proportion)	Session duration (seconds)
Heavy	**<50%**	0.47	0.45	82.36	0.42	435.40
	>50%	0.53	0.55	79.35	0.58	384.42
Medium	**<50%**	0.36	0.35	71.78	0.33	259.79
	>50%	0.64	0.65	65.41	0.67	217.07
Light	**<50%**	0.38	0.38	56.52	0.35	175.39
	>50%	0.62	0.62	52.08	0.65	163.34

Table 2: Descriptive statistics of daily phone usage grouped by battery level interval.

3.2 Features

The goal of this project is to identify distinct types of change in phone usage behavior by clustering smartphone users based on features that describe how smartphone users adapt their phone usage behavior to intervals of battery level. Therefore, all considered features represent a concept of phone usage behavior and quantify the extent to which a user adapts the respective concept given different intervals of battery level. This project focusses on the more general notions of phone usage behavior as discussed in the literature (e.g., mean application usage duration) and how a user adapts this behavior given two intervals of battery level (e.g., the difference in mean application usage duration between interval *x* and *y*).

Feature values have been computed for every participant in the dataset. The battery level variable has initially been binned using battery levels 1-50 and 51-100 as intervals. The observations of participants have been grouped by the defined intervals of battery level. For every participant, mean values of the respective concepts of phone usage behavior have been computed over both groups of battery level interval (e.g., a participant's mean application usage duration computed over observations belonging to interval 1-50 and interval 51-100 respectively). The computed mean value over a participant's observations belonging to the lower interval has been subtracted from the computed mean value over observations belonging to the higher interval. As a result, the features quantify to what extent a participant exhibits less or more of the respective concepts of phone usage behavior at the higher interval of battery level (e.g., if the user uses applications for shorter or longer periods of time when phone battery level is 51 percent or higher). A total of 104 feature values have been computed for every feature.

The features have been standardized to prevent the different magnitudes of feature values from influencing the cluster solutions. A robust scaler has been used to prevent outliers from biasing the scaling of values. The robust scaler subtracts the median value from feature values and scales the values using the 25th and 75th quantile. Standardization of feature values is solely done to prevent the different magnitudes of feature values from influencing the cluster solutions. Further sections of this report contain analyses of the original feature values.

Application usage duration. The application usage duration feature quantifies the difference in mean application usage time computed over battery level interval *x* and *y*. This feature indicates whether a user, on average, uses applications for longer or shorter periods of time at the higher interval of battery level. Figure 3 presents the distribution of feature values given the initial binning heuristic. With a mean feature value of -5.3 seconds (*SD* = 14.99), users tend to use applications for shorter periods of time at the higher interval of battery level.

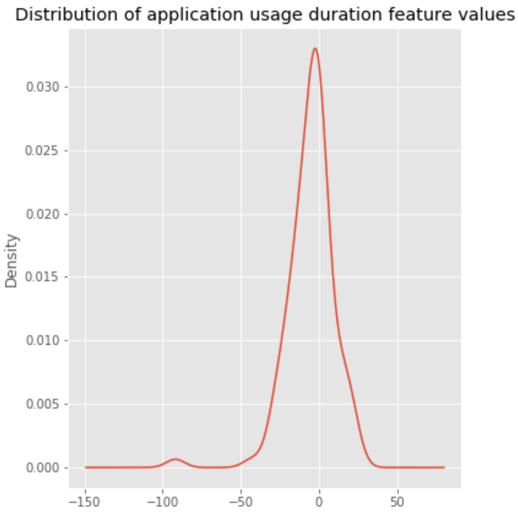

Figure 3: Distribution of the application usage duration feature values.

Number of apps used. The number of apps used feature quantifies the difference in the number of apps used computed over battery level interval *x* and *y*. This feature indicates whether a user uses less or more applications at the higher interval of battery level. On average, participants are likely to use more applications at the higher interval of battery level (*M* = 1,207.36, *SD* = 2,270.92). This might indicate that users are less likely to use applications at lower battery levels, or that users tend to keep their batteries charged at higher levels. Figure 4 presents the distribution of the number of apps used feature values.

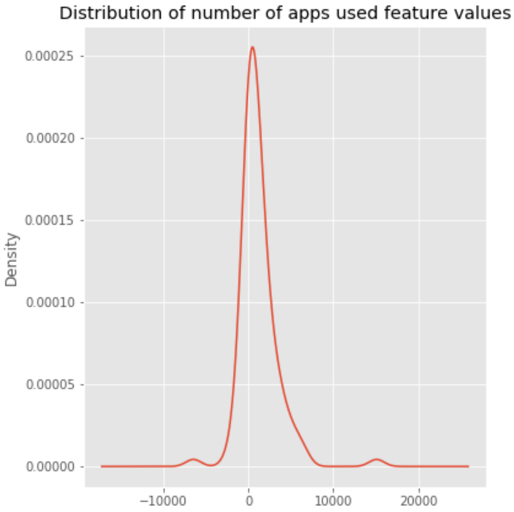

Figure 4: Distribution of the number of apps used feature values.

Session duration. The session duration feature quantifies the difference in mean sessions duration time computed over battery level *x* and *y*. This feature indicates whether a user's phone usage sessions are shorter or longer at the higher interval of battery level. On average, smartphone usage sessions last 38.19 seconds (*SD* = 81.87) shorter at the higher interval of battery level. Figure 5 presents the distribution of the session duration feature values.

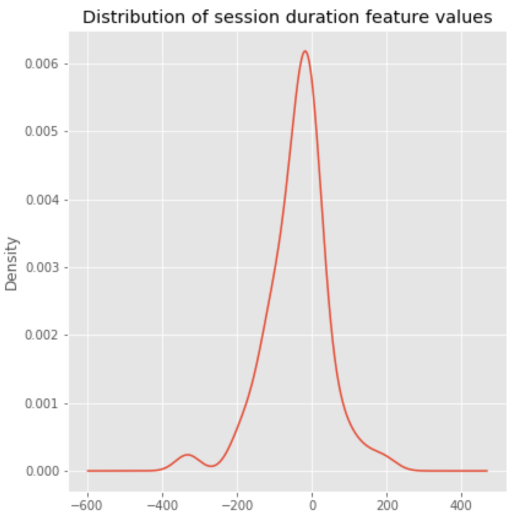

Figure 5: Distribution of the session duration feature values.

Number of sessions. The number of sessions feature quantifies the difference in the number of smartphone sessions computed over battery level *x* and *y*. This feature indicates whether a user has less or more smartphone usage sessions at the higher interval of battery level. The distribution of feature values is presented in figure 6. With a mean feature value of 451.58 (*SD* = 678.28), participants tend to initiate more smartphone usage session at the higher interval of battery level.

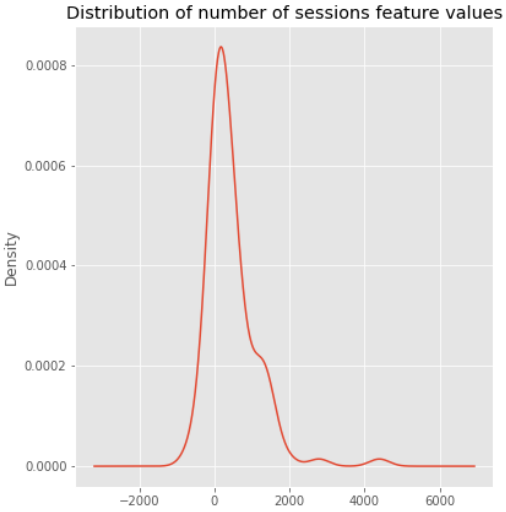

Figure 6: Distribution of the number of sessions feature values.

Apps per session. The apps per session feature quantifies the difference in mean number of applications used per sessions computed over interval *x* and *y*. This feature indicates whether a user, on average, uses less or more applications per session at the higher interval of battery level. Participants tend to use less applications per session at the higher interval of battery level (*M* = -0.31, *SD* = 1.26). Figure 7 presents the distribution of the apps per session feature values.

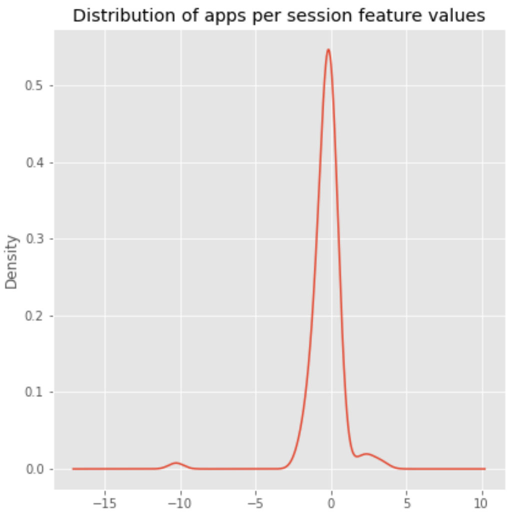

Figure 7: Distribution of the apps per session feature values.

Silence period. The silence period feature quantifies the difference in mean silence duration computed over interval *x* and *y*. Silence periods are periods during which a phone has not been used, and last from the end of the most recent session until the start of a new session. This feature indicates whether a user, on average, has longer or shorter periods of silence at the higher interval of battery level. Silence periods lasting longer than six hours have been considered as periods during which the participants sleep and were not included when computing the feature values. On average, silence periods last

133.34 seconds (*SD* = 592.43) shorter at the higher interval of battery level, indicating that people use their smartphones less frequently at lower battery levels. Figure 8 presents the distribution of the silence period feature values.

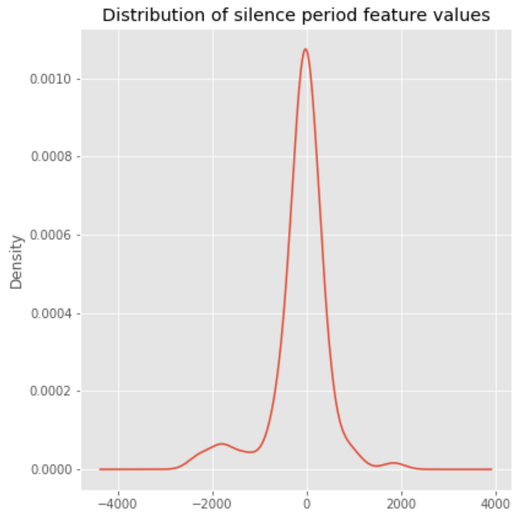

Figure 8: Distribution of the silence period feature values.

3.3 Method / Models

This project embodies an initial attempt at exploring how smartphone users adapt their phone usage behavior to the battery level of their phone. As formulated in the research questions, this research tries to evaluate to what extent distinct types of change in phone usage behavior can by derived by means of cluster analyses, and how stable the defined features and clustering solutions are when considering different intervals of battery level.

First, features have been created and capture the more general notions of phone usage behavior as discussed in the literature. The features have been combined into feature representations that attempt to conceptualize how smartphone users adapt their phone usage behavior to the battery level of their phone.

Second, the feature representations have been used to assess whether smartphone users show distinct types of change in phone usage behavior. The DBSCAN clustering algorithm has been used to account for small number of unique participants in the data set and to prevent noisy samples from seriously affecting the cluster results (Schubert, Sander, Ester, Kriegel & Xu, 2017). In addition, the density-based clustering algorithm is able to find clusters of arbitrary shapes, which can provide valuable information as they might indicate certain forms of change in smartphone usage behavior. The cluster solutions have been evaluated using the Density-Based Clustering Validation (DBCV) algorithm. DBCV is suited to assess density-based clustering models as it is able to evaluate globular as well as non-globular cluster solutions. The measure ranges from -1 to +1, where higher values imply better solutions (Moulavi, Jaskowiak, Campello, Zimek & Sander, 2014). Cluster solutions have been visualized to get a broad sense of how diverse or concise the different representations of change in phone usage behavior are given these global intervals. A total of 104 data points has been clustered. Cluster solutions have been established considering all possible feature representations. The DBSCAN clustering algorithm requires two hyperparameters, namely *minPts*, which quantifies the minimum amount of points required for a cluster to be formed, and *epsilon*, defining the maximum distance between points for them to be assigned to the same cluster (Schubert et al., 2017). According to Schubert et al. (2017), a *minPts* of four up to twice the number of dimensions suffices for most datasets, and *epsilon* should be set moderately small. In this project, *minPts* ranged from two up to twice the number of dimensions and *epsilon* was set to range for 0.1 up to the number of dimensions. The experiments have been conducted using the scikit-learn implementation of DBSCAN (Pedregosa et al., 2011).

Third, to gain a more gradual insight into how battery level affects phone usage behavior, multiple sets of features have been computed using different intervals of battery level. The initial cluster solutions have been established using features that have been computed using battery levels 1-50 and 51-100 as intervals. The solutions of this global approach might not hold when considering different, more narrow intervals of battery level. Different sets of features have been computed using the following pairings of battery level intervals: 0-20 & 20-40, 20-40 & 40-60, 40-60 & 60-80, 60-80 & 80-100. The stability of the features has been assessed by directly comparing the feature values of the sets of features. In addition, the newly computed sets of features have been clustered using the same feature representations and hyperparameters settings that have been used in the previous step. The cluster results have been compared to get insight in how stable the cluster solutions are across the different intervals of battery level.

4. Results

A clustering analysis has been conducted using the in section 3.2 defined features. The goal of the clustering analysis was to assess whether humans exhibit distinct types of phone usage adaptation to phone battery level and how different intervals of battery level affect this adaptation process.

Table 3 presents the top ten cluster results in terms of clustering quality according to the DBCV measure. With a minimum value of 0.709, all cluster results are of considerable quality. Every cluster solution consists of only one cluster. The results indicate that users do not show clear distinct types of smartphone usage adaptation as defined by the feature representations.

	Hyperparameters		**Evaluation metrics**		
Representation	**minPts**	**Epsilon**	**DBCV**	**Noise**	**Clusters**
1. f1, f3, f5	3	3.7	0.758	2	1
2. f2, f3, f4, f5	3	4.9	0.744	2	1
3. f2, f4, f5	3	3.7	0.739	2	1
4. f2, f3, f4	3	2.5	0.738	2	1
5. f1, f2	2	2.1	0.725	3	1
6. f2, f4	2	2.5	0.725	2	1
7. f2, f3	2	2.1	0.719	2	1
8. f5, f6	2	2.5	0.718	2	1
9. f1, f3, f4, f5	4	4.9	0.717	2	1
10. f1, f4, f5	2	3.3	0.709	2	1

Table 3: Top 10 cluster results. F1) App usage duration, F2) Number of apps used, F3) Session duration, F4) Number of sessions, F5) Apps per session, F6) Silence period.

Cluster results have been plotted to gain insight into how the different feature spaces look like. The most interesting feature spaces have been discussed in this section. The remaining plots can be found in the appendices. Every cluster result has been plotted twice, one plot visualizing the cluster labels and noise while the other plot visualizes the distribution of user classes.

The feature representation of cluster result 1 is visualized in figure 9. There appears to be one dense area of cluster points surrounded by additional cluster points and two clear outliers. There is little to no indication suggesting the forming of additional clusters. Looking at the class distribution, there is no clear sign that user classes show distinct behavior in terms of the used features.

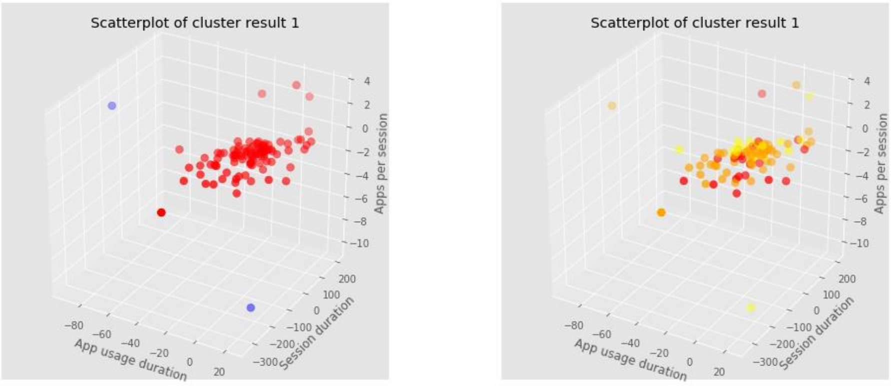

Figure 9: Scatterplots of cluster result 1. In the left plot, red dots indicate cluster labels and blue dots represent observations flagged as noise. In the right plot, colors represent the user types. Heavy, medium and light users are colored red, orange and yellow respectively.

Cluster result 4 is visualized in figure 10. Cluster result 4 is approximately linearly shaped and consists of one dense area of cluster points surrounded by a few distant cluster points and noise. Similar to cluster result 1, there is no indication suggesting the forming of additional clusters.

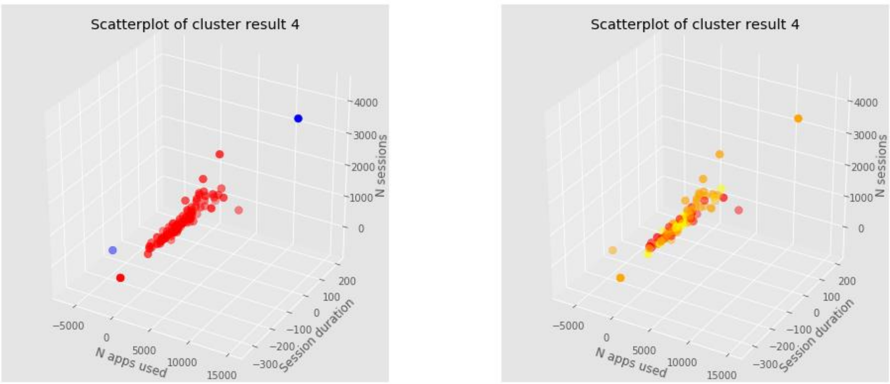

Figure 10: Scatterplots of cluster result 4. In the left plot, red dots indicate cluster labels and blue dots represent observations flagged as noise. In the right plot, colors represent the user types. Heavy, medium and light users are colored red, orange and yellow respectively.

Upon closer inspection of the used features in cluster result 4, we find what appears to be a linear relationship between the 'Number of sessions' and 'Number of apps used' feature (figure 11). This relation logically makes sense as users that use more apps are likely to use more sessions as well.

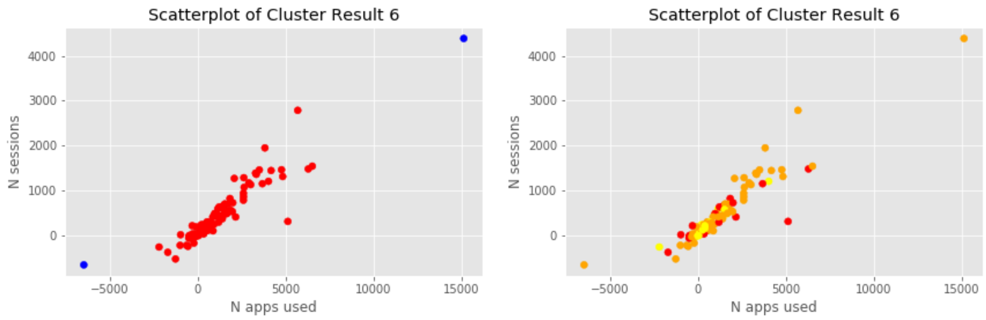

Figure 11: Scatterplots of cluster result 6. In the left plot, red dots indicate cluster labels and blue dots represent observations flagged as noise. In the right plot, colors represent the user types. Heavy, medium and light users are colored red, orange and yellow respectively.

Figure 12 visualizes the feature space of cluster solution 8. Contrary to the other cluster solutions, this result is less dense and suggests the forming of a second cluster. However, the DBSCAN algorithm still considered this to be one cluster because all points lie relatively close to each other. The distribution of user classes implies that light users use a consistent amount of applications but show much more inconsistent behavior in terms of silence periods.

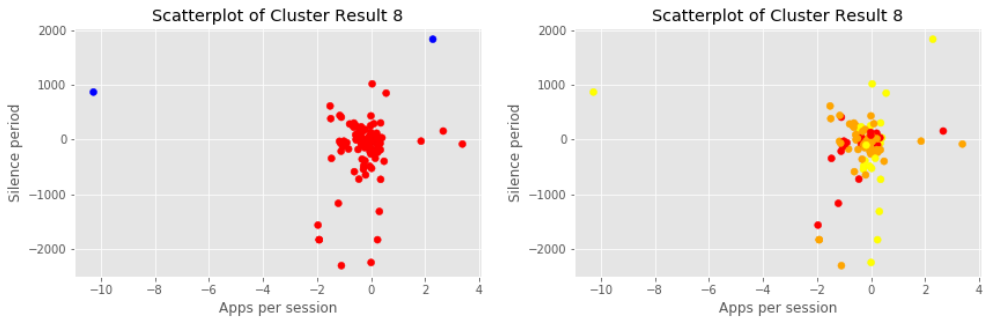

Figure 12: Scatterplots of cluster result 8. In the left plot, red dots indicate cluster labels and blue dots represent observations flagged as noise. In the right plot, colors represent the user types. Heavy, medium and light users are colored red, orange and yellow respectively.

The discussed cluster results are based on features that captured change in smartphone usage behavior using battery levels 1-50 and 51-100 as intervals. Table 4 presents mean feature values computed over different intervals of battery level. We find that the used intervals (0-50, 51-100) capture the change in features 1 and 5 relatively well as the mean values are rather stable across the different intervals of battery level. However, features 2, 3, 4 and 6 clearly show a much more gradual increase or decrease when computed over different intervals.

	Interval 0-20 / 20-40	Interval 20-40 / 40-60	Interval 40-60 / 60-80	Interval 60-80 / 80-100	Interval 0-50 / 50-100
Feature 1	0.16	-1.42	-5.35	4.87	-5.3
Feature 2	388.01	218.5	121.36	121.46	1207.36
Feature 3	23.39	-28.19	-106.4	-182.24	-38.2
Feature 4	126.32	88.41	54.66	26.43	451.58
Feature 5	-0.21	-0.17	-0.18	0.07	-0.31
Feature 6	270.81	-146.22	-118.1	-65.57	-133.34

Table 4: Average feature values computed over different intervals of battery level. F1) App usage duration, F2) Number of apps used, F3) Session duration, F4) Number of sessions, F5) Apps per session, F6) Silence period.

Using the same feature representations and hyperparameters settings, the top ten clustering solutions have been reevaluated using the different sets of feature values. Table 5 presents the results. It indicates that, despite configurations being exactly similar, the cluster solutions are quite unstable across the different intervals of battery level. The DBCV measure resulted in nan values when all points have been flagged as noise.

	DBCV Interval 0-20 / 20-40	DBCV Interval 20-40 / 40-60	DBCV Interval 40-60 / 60-80	DBCV Interval 60-80 / 80-100	DBCV Interval 0-50 / 50-100
f1, f3, f5	0.414	0.353	-0.065	0.627	0.758
f2, f3, f4, f5	0.273	0.482	0.459	0.629	0.744
f2, f4, f5	0.627	0.352	0.483	0.576	0.739
f2, f3, f4	-0.226	nan	0.256	0.661	0.738
f1, f2	0.56	-0.543	0.233	0.029	0.725
f2, f4	-0.027	nan	0.492	0.501	0.725
f2, f3	-0.29	0.583	0.396	0.686	0.719
f5, f6	0.558	0.48	-0.151	-0.345	0.718
f1, f3, f4, f5	0.564	0.389	-0.18	0.615	0.717
f1, f4, f5	0.481	0.246	0.013	0.541	0.709

Table 5: Cluster results using the newly computed feature values. F1) App usage duration, F2) Number of apps used, F3) Session duration, F4) Number of sessions, F5) Apps per session, F6) Silence period.

5. Discussion

The goal of this study was to ascertain how smartphone users adapt their phone usage behavior to the battery level of their phone. The project attempted to evaluate to what extent distinct types of change in phone usage behavior can by derived by means of cluster analyses, and how stable the defined features and clustering solutions are when considering different intervals of battery level.

This study derived six features from smartphone log data which all represented a concept of phone usage behavior. All features represented a concept of phone usage behavior (e.g., mean application usage duration) and quantified to what extent a user adapts the respective concept of phone usage behavior given two intervals of battery level (e.g., the difference in mean application duration between interval *x* and *y*). The features have been combined into different feature representations that attempt to conceptualize how smartphone users adapt their phone usage behavior to the battery level of their phone.

Using the defined feature representations as input, the DBSCAN clustering algorithm was used to assess whether smartphone users show distinct patterns of phone usage adaptation to phone battery level. We found that the top 10 cluster solutions in terms of the Density-Based Clustering Validation measure (DBCV) all consist of only one cluster and a small amount of noise. These results strongly indicate that smartphone users do not show distinct patterns of change in phone usage behavior as defined by the evaluated feature representations.

With a minimum DBCV measure of 0.709, the established cluster solutions are of considerable quality. An important thing to keep in mind is that the DBCV measure, similar to measures like Silhouette, indicates how similar points are to their own cluster as compared to other clusters. However, given that all analyzed solutions only consist of one cluster, the DBCV measures indicated how similar points in the established clusters are compared to the points that were assigned as noise. It is therefore questionable how useful the metric is in this context and whether or not a density-based clustering algorithm is a meaningful method to analyze the data of interest.

Visualizing the cluster solutions in feature space verified that the established clusters were of substantial density. The visualizations provided little to no indication that hinted the forming of additional clusters. Instead, they suggest that phone usage adaptation fluctuates more gradually between users. Additionally, the visualizations indicated that the ‘Number of sessions’ and ‘Number of apps used’ features are linearly related, suggesting that users who use more applications simultaneously initiate more smartphone sessions. This relation logically makes sense.

The cluster results strongly indicated that smartphone users do not show distinct patterns of change in phone usage behavior as defined by the evaluated feature representations. Instead, visualizations of the cluster solutions imply that users distinct themselves from others by marginally adapting more or less extreme to different features. However, the feature values have been computed using battery levels 1-50 and 51-100 as intervals. Further analysis of the used feature values suggested that these intervals are too broad and indicated that the cluster solutions do not accurately capture the natural flow of phone usage adaptation. By binning the battery variable in five equally sized intervals, new sets of feature values have been computed considering all pairs of neighboring intervals. By analyzing the mean feature values, it was found that four of the six features exponentially increase or decrease as battery level declines. This information has not been captured by the features that have been used for clustering.

Finally, the stability of the cluster solutions has been tested using the newly computed feature values. Using the same feature representations and hyperparameter settings that have been used with the top ten cluster solutions, we assessed how the quality of the cluster solutions changed. All new cluster solutions were of less quality than the established ones. These results suggest that phone usage adaptation is best to be captured using more narrowly defined intervals. A plausible explanation could be that smartphone users distinct themselves from others by adapting their phone usage behavior more or less aggressively at different intervals of battery level. However, we must keep in mind that the features are computed using less data as the amount of battery intervals increase. This drawback might have partially caused the decrease in cluster quality.

6. Conclusion

This study attempted to answer the following research questions:

- To what extent can we distinct different patterns of phone usage adaptation to phone battery level by clustering smartphone users based on features that describe how smartphone users adapt their phone usage behavior to intervals of battery level?
- Are the defined features and cluster solutions stable when we evaluate how smartphone users adapt their phone usage behavior considering different intervals of battery level?

The cluster results indicated that smartphone users do not show distinct patterns of change in phone usage behavior. The analyzed cluster results all consist of only one cluster and a small amount of noise. The cluster results did not show strong indications hinting the forming of additional clusters. Instead, the they suggest that phone usage adaptation fluctuates more gradually between users.

However, the cluster results have been established using features that captured the change in phone usage behavior using battery levels 1-50 and 51-100 as intervals. Further analysis of the used features strongly indicated that this approach is too global: feature values increased or decreased exponentially as battery level decreases. These findings indicate that phone usage behavior is best captured using multiple intervals of battery level.

The stability of the cluster results has been tested using features computed over more narrow intervals of battery level. All new cluster solutions were of less quality than the established ones. These results further suggest that the approach using battery levels 1-50 and 51-100 as intervals is too broad. A plausible explanation could be that smartphone users distinct themselves from others by adapting their phone usage behavior more or less aggressively at different intervals of battery level.

This work has several implications for related fields of research. For example, it found that the used concepts of phone usage behavior exponentially vary as battery level decreases which partially explains the high variability of phone usage in general. Moreover, the results strongly suggest that change in phone usage behavior as a function of battery level cannot be accurately captured on a global scale. Instead, we should look closely to how smartphone users adapt their phone usage behavior on a more continuous scale. This might prove to be very relevant for works that attempt to predict battery lifetime.

Future research could expand on this work by evaluating the concept of change in phone usage behavior as a function of battery level on a more continuous scale. This can be achieved through expansion of the used feature representations by computing additional sets of features using more narrowly defined intervals of battery level. Future work could also include more concepts of change in phone usage behavior such as application micro-usage (Ferreira et al., 2014) and the use of various types of applications.

References

Ahmad, R. W., Gani, A., Hamid, S. H. A., Shojafar, M., Ahmed, A. I. A., Madani, S. A., ... & Rodrigues, J. J. (2017). A survey on energy estimation and power modeling schemes for smartphone applications. *International Journal of Communication Systems*, *30*(11), e3234.

Anjomshoa, F., & Kantarci, B. (2018). Sober-mcs: Sociability-oriented and battery efficient recruitment for mobile crowd-sensing. *Sensors*, *18*(5), 1593.

Banerjee, N., Rahmati, A., Corner, M. D., Rollins, S., & Zhong, L. (2007, September). Users and batteries: interactions and adaptive energy management in mobile systems. In *International conference on ubiquitous computing* (pp. 217-234). Springer, Berlin, Heidelberg.

Brown, B., McGregor, M., & McMillan, D. (2014, September). 100 days of iPhone use: understanding the details of mobile device use. In *Proceedings of the 16th international conference on Human-computer interaction with mobile devices & services* (pp. 223-232). ACM.

Chantrapornchai, C., & Nusawat, P. (2016). Two Machine Learning Models for Mobile Phone Battery Discharge Rate Prediction Based on Usage Patterns. *Journal of Information Processing Systems*, *12*(3).

Church, K., Ferreira, D., Banovic, N., & Lyons, K. (2015, August). Understanding the challenges of mobile phone usage data. In *Proceedings of the 17th International Conference on Human-Computer Interaction with Mobile Devices and Services* (pp. 504-514). ACM.

Datta, S. K., Bonnet, C., & Nikaein, N. (2014, May). Personalized power saving profiles generation analyzing smart device usage patterns. In *2014 7th IFIP Wireless and Mobile Networking Conference (WMNC)* (pp. 1-8). IEEE.

Draa, I. C., Niar, S., Tayeb, J., Grislin, E., & Desertot, M. (2017). Sensing user context and habits for run-time energy optimization. *EURASIP Journal on Embedded Systems*, *2017*(1), 4.

Draa, I. C., Niar, S., Grislin-Le Strugeon, E., Biglari-Abhari, M., & Tayeb, J. (2019). ENOrMOUS: ENergy Optimization for MObile plateform using User needS. *Journal of Systems Architecture*, *97*, 320-334.

Deng, D. (2015). Li-ion batteries: basics, progress, and challenges. *Energy Science & Engineering*, *3*(5), 385-418.

Falaki, H., Mahajan, R., Kandula, S., Lymberopoulos, D., Govindan, R., & Estrin, D. (2010, June). Diversity in smartphone usage. In *Proceedings of the 8th international conference on Mobile systems, applications, and services* (pp. 179-194). ACM.

Ferreira, D., Dey, A. K., & Kostakos, V. (2011, June). Understanding human-smartphone concerns: a study of battery life. In *International Conference on Pervasive Computing* (pp. 19-33). Springer, Berlin, Heidelberg.

Ferreira, D., Goncalves, J., Kostakos, V., Barkhuus, L., & Dey, A. K. (2014, September). Contextual experience sampling of mobile application micro-usage. In *Proceedings of the 16th international conference on Human-computer interaction with mobile devices & services* (pp. 91-100). ACM.

Guo, Y., Wang, C., & Chen, X. (2017). Understanding application-battery interactions on smartphones: A large-scale empirical study. *IEEE Access*, *5*, 13387-13400.

Hendrickson, A. T., Aalbers, G., & Vanden Abeele, M. (under review)

Kekolahti, P., Kilkki, K., Hämmäinen, H., & Riikonen, A. (2016). Features as predictors of phone popularity: an analysis of trends and structural breaks. *Telematics and Informatics*, *33*(4), 973-989.

Kim, D., Chon, Y., Jung, W., Kim, Y., & Cha, H. (2016). Accurate prediction of available battery time for mobile applications. *ACM Transactions on Embedded Computing Systems (TECS)*, *15*(3), 48.

Li, H., Liu, X., & Mei, Q. (2018). Predicting smartphone battery life based on comprehensive and real-time usage data. *arXiv preprint arXiv:1801.04069*.

Mirsky, Y., Shabtai, A., Rokach, L., Shapira, B., & Elovici, Y. (2016, October). Sherlock vs moriarty: A smartphone dataset for cybersecurity research. In *Proceedings of the 2016 ACM workshop on Artificial intelligence and security* (pp. 1-12). ACM.

Moulavi, D., Jaskowiak, P. A., Campello, R. J., Zimek, A., & Sander, J. (2014, April). Density-based clustering validation. In *Proceedings of the 2014 SIAM International Conference on Data Mining* (pp. 839-847). Society for Industrial and Applied Mathematics.

Oliver, E. A., & Keshav, S. (2011, September). An empirical approach to smartphone energy level prediction. In *Proceedings of the 13th international conference on Ubiquitous computing* (pp. 345-354). ACM.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., ... & Vanderplas, J. (2011). Scikit-learn: Machine learning in Python. *Journal of machine learning research*, *12*(Oct), 2825-2830.

Schubert, E., Sander, J., Ester, M., Kriegel, H. P., & Xu, X. (2017). DBSCAN revisited, revisited: why and how you should (still) use DBSCAN. *ACM Transactions on Database Systems (TODS)*, *42*(3), 19.

Yurur, O., Liu, C., & Moreno, W. (2015). Modeling battery behavior on sensory operations for context-aware smartphone sensing. *Sensors*, *15*(6), 12323-12341.

Zhao, S., Ramos, J., Tao, J., Jiang, Z., Li, S., Wu, Z., ... & Dey, A. K. (2016, September). Discovering different kinds of smartphone users through their application usage behaviors. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing* (pp. 498-509). ACM.

Appendix A: Visualizations of cluster results

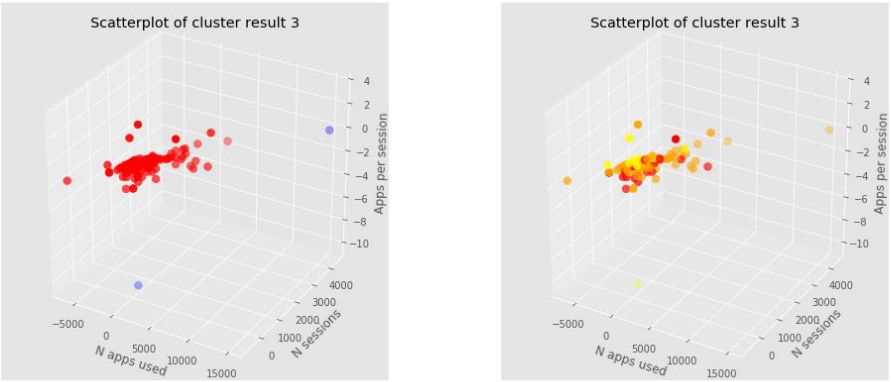

Figure 13: Scatterplots of cluster result 3. In the left plot, red dots indicate cluster labels and blue dots represent observations flagged as noise. In the right plot, colors represent the user types. Heavy, medium and light users are colored red, orange and yellow respectively.

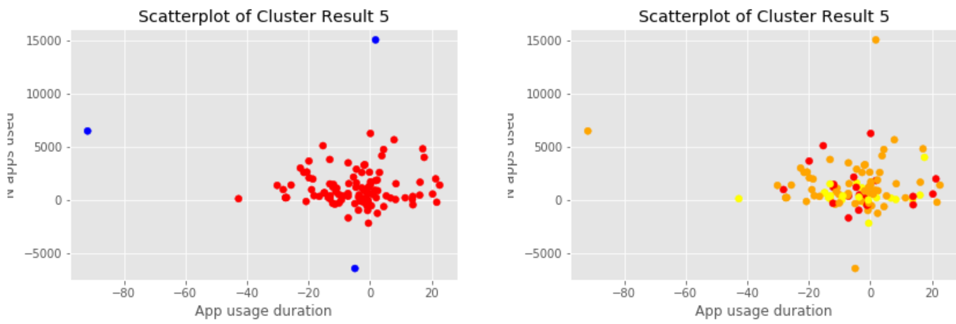

Figure 14: Scatterplots of cluster result 5. In the left plot, red dots indicate cluster labels and blue dots represent observations flagged as noise. In the right plot, colors represent the user types. Heavy, medium and light users are colored red, orange and yellow respectively.

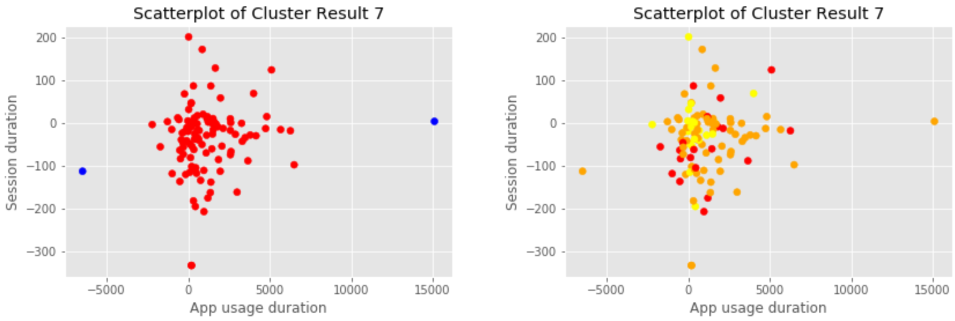

Figure 15: Scatterplots of cluster result 7. In the left plot, red dots indicate cluster labels and blue dots represent observations flagged as noise. In the right plot, colors represent the user types. Heavy, medium and light users are colored red, orange and yellow respectively.

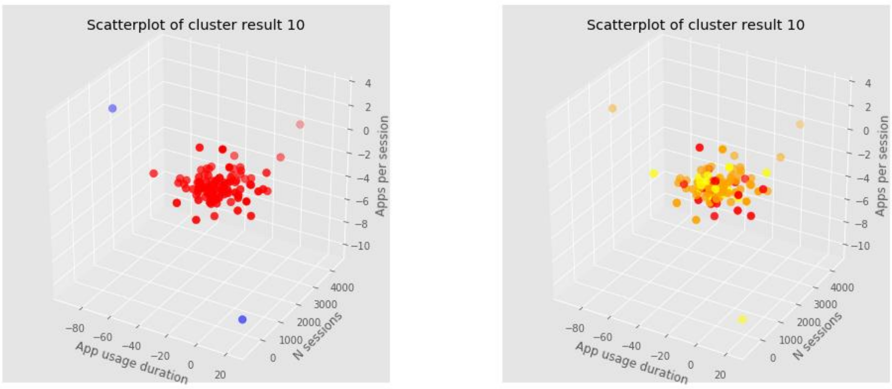

Figure 16: Scatterplots of cluster result 10. In the left plot, red dots indicate cluster labels and blue dots represent observations flagged as noise. In the right plot, colors represent the user types. Heavy, medium and light users are colored red, orange and yellow respectively.