

Exploring the effects of below-freezing temperatures on smartphone usage

Aku Visuri^a, Jonatan Hamberg^b, Ella Peltonen^{a,*}

^a Center for Ubiquitous Computing, University of Oulu, P.O. Box 4500, FI-90014, Finland

^b Department of Computer Science, University of Helsinki, P.O. Box 64, FI-00014, Finland

ARTICLE INFO

Article history:

Received 24 February 2021

Received in revised form 8 November 2021

Accepted 11 November 2021

Available online 27 November 2021

Keywords:

Smartphones

Temperatures

Mobile sensing

ABSTRACT

While the use of smartphones in extreme temperatures does not necessarily occur every day nor in all parts of the world, numerous use cases can be highlighted where the use of smartphones in cold temperatures is mandatory. Modern smartphones are designed to function in a wide range of temperatures, but when exposed to extreme cold temperatures the performance and reliability can significantly suffer. This paper presents a controlled laboratory experiment, using a clinical cold chamber to expose seven smartphone models to both medium cold (0 °C to −20 °C) and extreme cold (−30 °C) environments. The results showcase the smartphones' sensing software's lack of awareness of the cold environment, as well as reliability issues in the form of device crashes across the whole range of tested devices. We present a strategy for implementing monitoring application designs to both appropriately sense the effect of cold environments, as well as predicting device shutdowns in extreme cold.

© 2021 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Extreme cold temperatures are common in the Northern hemisphere, including Scandinavia, North America, and North Asia. In many parts of the world, diurnal and seasonal cycles can produce cold temperatures down to negative degrees Celsius. People working, commuting or otherwise spending time in outdoor activities cannot always avoid exposing their smartphones to below freezing temperatures. To motivate the need for our research, we ran an online survey for people regularly using smartphones in below freezing conditions (N = 129, 58% female). Out of these answers coming all around the Nordic hemisphere, 75% reported regularly using their smartphones in such conditions for commuting apps, such as maps, and 87% for making calls or sending messages. Other reported use cases were listening to music, playing games when waiting for something, such as public transportation, and taking pictures of snowy landscapes or northern lights.

Because the situation is commonplace for Nordics, cold-related malfunctions and usage failures in smartphones have been widely reported. Based on our online survey, most commonly met malfunctions were unexpected drops in battery level (78% out of study participants using smartphones in cold) and unexpected app or operating system crashes (37%). Even though, to date, these general observations have not been scientifically verified, similar problems have been previously noted, e.g., in the context of electric and hybrid vehicles [1].

While the effects of extreme environments have been previously studied from the perspective of smartphone usability and interactions [2], it is unclear how long the devices can withstand cold conditions and handle transitions from extreme environments (e.g. from indoors to outdoors and vice versa). Previous studies have shown that the battery life of

* Corresponding author.

E-mail addresses: aku.visuri@oulu.fi (A. Visuri), jonatan.hamberg@helsinki.fi (J. Hamberg), ella.peltonen@oulu.fi (E. Peltonen).

smartphones is affected by the device temperature [3], but mainly in association with their use rather than the ambient temperature. Because prolonged activities in subzero conditions can pose an increased health risk, and unexpected subzero periods can expose us to severe health condition and even increase mortality [4], it is extremely important that safety and emergency features of smartphones continue working in all situations. In addition, malfunctions in smartphones during sudden arctic outbreaks and other climatic fluctuations can further exacerbate chaos in transportation and aggravate the already unpleasant conditions.

Batteries in modern smartphones are equipped with software-enabled sensors that periodically report the battery state, including but not limited to its temperature and health, to the operating system. Using this information mobile operating systems should, ideally, be able to react to battery malfunctions and other special conditions, such as cold temperatures below the operating range of the device. However, since implementation details of the battery health feature are left to the original equipment manufacturer, not all devices accurately report the diagnostic. In the wild, smartphones are rarely observed warning the user about potential instability before it is too late. Fortunately, both iOS and Android provide access to battery temperature readings through a specialised operating system API, allowing other applications to observe and utilise the data.

Based on our online survey, people using smartphones in cold environments were not aware of precautions they could take to keep their smartphone operating longer period of time. Some reported carrying it close to body to keep the phone relatively warm (61% out of 129 respondents) but some took no precautions at all but used the phone as usual (35%). Another important highlight of this questionnaire is that people indeed use their smartphones in cold and they do it daily basis during the winter months. Even if the issue might seem to be small or avoidable for those not living in the Arctic hemisphere, locals consider smartphones as crucial part of their everyday life and temperature-driven problems a true nuisance.

In this paper, we explore the effects of extreme cold on smartphone performance, measure the accuracy of reported battery health readings, and conclude that in-built diagnostics in mobile operating systems cannot always sufficiently react to cold-temperature interference or warn the user before instability. For this purpose, we thoroughly assessed the performance of six commercial Android and one iOS smartphone in cold (0 °C, -10 °C) and extreme cold (-20 and -30 °C) temperatures using two controlled cold chambers. We collect the battery data (such as battery level, temperature, and voltage) for the full discharging loop, from 100% battery level to termination point, with different CPU loads, and compare the results in different cold temperatures.

In our analysis, we focus the following aspects, or research questions, of the smartphone usage in extremely cold: First, we investigate what influences the change in room temperature has on the on-device battery levels. Second, we study what malfunctions appear when smartphones are exposed for extreme cold. Lastly, we explore if a device shut down due to exposure to extreme cold temperature can cause long-term damage to the battery that could result in future malfunctions, by modelling battery health awareness and crash prevention. The contributions of this paper can be summarised as the following:

- We showcase an experimental setup for exploring smartphones to cold temperature in a controlled environment, where we can see that cold ambient temperature, indeed, has statistically significant impact on battery usage.
- We present accurate prediction models for reporting the battery health status and foreseeing a likely crash event, showcasing there are over 27% critical exposures unobserved by the operating system.
- We propose a technique to alert the user of possible device instability due to cold weather and discuss the potential preventive measures for keeping the device functional in extreme cold temperatures.

2. Related work

This paper focuses on the detection and prevention of smartphone anomalies caused by extremely cold ambient temperatures. While modern smartphones can operate in a wide range of temperatures, both overheating and extreme cold can cause problems with device stability and battery life [5]. In many habitats with cold climates, including but not limited to Northern Europe, Canada, and high-altitude locations, users frequently interact with their devices in outdoor settings. This exposure to cold is usually limited but may be prolonged due to occupational requirements or recreational outdoor activities [6,7]. Since these conditions are known to deteriorate work performance [8] and impair interaction with devices [9], problems with device stability can further amplify the negative effects of cold temperature. While the pervasive community has taken steps to go “beyond the desktop” [10], only a handful of works have explored the use of devices in non-conventional environments (summarised in Table 1).

Temperature Measurements. Batteries in modern smartphones are equipped with temperature sensors in order to protect them from damage. Overeem et al. [11] and Droste et al. [12] show that readings from these sensors correlate with daily mean air temperatures (10–30 °C). These can be considered normal operating temperatures for common smartphone usage, and the results leave it unclear whether similar outcome could be expected in any colder temperatures (minus degrees Celsius). Studies of smartphone behaviour in cold has only been briefly considered previously [9], focusing on user interaction of the smartphone screen in a cold environment.

Internal Heat Production. Smartphones and other hand-held electronic devices by themselves are considered as heat sources [11]. Xie et al. [13] mention especially CPU and battery as the major heat-generating components. Gurrum et al. [14] ran a study of temperature measurements for different smartphone hardware components, such as battery,

Table 1
Summary of the previous contributions in comparison to this paper.

Reference	Data	Findings
Vlahinos et al. [5], 2002	Electric and hybrid vehicle batteries	Batteries require heating in cold environments to function appropriately.
Overeem et al. [11], 2013, and Droste et al. [12], 2017	Smartphone heat sensor and environmental temperature data	In warm (10–30C) temperatures the smartphone sensor correlates to surrounding temperature.
Zhang et al. [25], 2010, and Kjaergaard [22], 2010	Generic smartphone energy profiles	Power models and profiles can enable developers to create more power-efficient applications.
Pathak et al. [24], 2012, and Liu et al. [26], 2013	Smartphone application energy profiles	Identified numerous energy inefficiency problems in smartphone applications.
Ma et al. [27], 2013, and Oliner et al. [28], 2013	General smartphone battery drain anomalies and causes	Both applications (eDoctor and Carat) managed to diagnose and prevent anomalous battery drain.
This paper	Smartphone battery draining profiles	Focus on malfunctions in extreme cold environments.

display, and skin plates. Therminator [15] simulate relationship between temperature of hardware components and skin layers of smartphones. To summarise, hardware components, such as CPU, GPU, and battery, can produce a temperature varying from 30 to 50 Celsius degrees in normal room conditions. These temperature levels have also been confirmed through the battery temperature sensor of smartphones when studying the effect on battery life [3]. When considering a smartphone as a thermal sensor, one must analyse and consider the heat generated by the device itself in addition to the ambient temperature.

Battery Performance. Studies of Lithium-ion batteries and temperature usually focus on cooling the battery-dependent devices [16,17]. This is, indeed, a more common use case: usually we want to lower the battery temperature for better performance (in normal operational ambient temperatures), but extreme conditions can cause unwanted behaviour in batteries. Specifically, previous work has not systematically investigated how device battery life and overall performance are impacted by subzero ambient temperatures. Vlahinos et al. [5] show that in extreme cold scenarios, batteries can sometimes require heating to function. Wang et al. [18] propose a solution of self-heating batteries focusing on vehicles and high-altitude drones, but it remains unclear if this type of hardware component could fit into a hand-held device.

Background Temperature Control. In this paper we conduct an experiment using a specially designed medical-purpose climatic chamber. This chamber allows us to control the environment temperature and humidity for modelling and experiments. This approach has been previously used, for example, when studying Bluetooth environmental sensors [19], batteries [20], and user interaction in cold environments [9], and in numerous medical experiments related to short- and long-term exposure to cold climates [21].

Energy Consumption. Several prior studies have analysed general mobile device energy profiles [22–25], or application specific energy use [24,26–28]. While these works cover a wide range of different modelling contexts, ambient temperature is absent from use as a predictor or an outcome variable. Battery temperature was shown to have a large impact on energy consumption [3], but as the battery temperature has a close relationship with other physical components (and their use), the temperature in normal conditions could simply be a by-product of e.g., high CPU load. This paper explores the variability of ambient temperature and how subzero temperatures influence smartphone battery performance. As such, this is the first work to perform such an experiment in a laboratory setting.

3. Method and experimental setup

3.1. Physical environment

Our experiments were conducted in two climatic cold chambers at the Oulu Regional Institute of Occupational Health, Oulu, Finland. Both chambers were air convection cooled to a controller ambient temperature with an accuracy of ± 2 °C. The first chamber was used for medium cold conditions ranging from -20 °C to 0 °C, while the second one enabled more extreme temperatures down to -30 °C. Both chambers had wind velocity below 0.1 m/s and a humidity of 30%–35% to simulate the dry characteristics of cold climates.

In the cold chambers, devices were seated on a wooden table in separate polystyrene foam bases to prevent heat dissipation (left in Fig. 1). To maximise the area of exposure while keeping the devices stationary, all smartphones were docked upright in 1–2 cm holes in the base material. When docked, all parts of the device case protecting core components (e.g. CPU, battery) were still directly exposed to the cold ambient temperature. The chambers and target devices were remotely monitored during measurements in order to avoid discomfort and thermal inference from external sources.




Fig. 1. From left to right: cold chamber environment, room control software, and a phone after measurements. Note the frost on the display.




Fig. 2. Experimental setup.

Table 2

Devices used in the experiment. Nominal operating temperature for all the devices is 0–35 °C.

Manufacturer and model	Operating system	Battery capacity
Apple A1457 (iPhone 5)	iOS 9.3.5	1440 mAh
Samsung G900F (S5)	Android 6.0.1	2800 mAh
Samsung G920F (S6)	Android 7.0	2550 mAh
Samsung i9100 (S2)	Android 4.1.2	1650 mAh
Samsung i9250 (Nexus)	Android 4.0.1	1750 mAh
Samsung i9300 (S3)	Android 4.3	2100 mAh
Samsung i9506 (S4)	Android 5.0.1	2600 mAh

Since both cold chambers were thermally insulated, the target devices had very poor or completely nonexistent cellular connectivity in them. The collected data is stored on the smartphone and synced to a server periodically. In order to remotely monitor the smartphones, we installed a wireless 2.4 GHz Wi-Fi access point (WAP) which had a signal strong enough to penetrate the cold chamber wall when placed in an adjacent room (see Fig. 2). In the WLAN provided by the WAP, we also setup a local server for monitoring and logging the state of devices. It was necessary for the server to have a display, as it was used to monitor the exit conditions of the experiments (e.g. device shutdown). The Wi-Fi connectivity was also used to ensure data backups in case the smartphones crash in extreme conditions and cease to function, hindering the stored data inaccessible.

3.2. Experiment procedure

Experiment smartphones. Our experiment was performed on smartphones listed in Table 2, six Samsung devices running Android and one iPhone with iOS. At the time of the experiments, used devices were two to four years old and served their active lifetime mainly as educational tools in programming courses. This means they were definitely less used than regular smartphones in personal everyday use, but old enough to “retire” without large financial costs even if something

Table 3
Experiment setup: cases to be run for the devices (see Table 2).

CPU 100%	CPU 100%	CPU 100%	CPU 100%
0 °C	All devices	All devices	All devices
−10 °C	All devices	All devices	All devices
−20 °C	All devices	All devices	All devices
−30 °C	All but S5, S6	All but S5, S6	All but S5, S6

Table 4
System settings preset.

Setting	State
Screen	On
Screen timeout	Indefinite
Brightness	Maximum
WiFi	On
Mobile network	Off
Location	Off
Bluetooth	Off

went wrong with the cold exposure. Both Samsung and Apple guarantee nominal operating temperature of 0–35 °C for their smartphones, which means our experiments included a risk of harming the electronics.

Data collection applications. We periodically captured internal subsystem variables using specifically tailored logging applications developed by us: two of which dedicated for logging on Android and iOS respectively and one for generating artificial CPU usage on Android. Each logging application was implemented as a background service which queried and stored its results to the internal storage of the device in 1 s intervals using NTP timestamps independent from the local system clock, allowing for a robust synchronisation of the device readings. Both Android and iOS apps got the readings through their regular APIs.

Experiment setup. The conducted experiments are summarised in Table 3. All the experiments were run for the test smartphones in 0, −10, and −20 °C temperatures, and with three different CPU loads (artificial loads 50% and 100%, and without artificial load), discharging from the full battery until the termination of the experiment. Two of the newest device models (S5, S6 at the time) were not taken into the coldest (−30 °C) temperature per request of the administrator, as the initial hypothesis was that this temperature could permanently damage the devices. The older devices (at the time) were also measured in −30 °C.

Prior the experiments. Prior to starting the measurements, we conducted a preparation phase for the target devices where they were warmed up to the ambient room temperature, charged to 100% battery and fixed to a preset described in Table 4 to unify the initial system states. In all experiments, the screen brightness was set to maximum and its timeout disabled to simulate outdoor conditions in the daylight. All non-system applications were closed, and automatic updates disabled in order to avoid unexpected spikes in energy consumption. Other networking than Wi-Fi, such as Bluetooth and GPS, were disabled as well. Few moments before starting the measurement, artificial CPU load was set regarding the case in question.

Termination of the experiments. The aim of the study was keep the devices in the cold chamber for entirety of their battery cycle, including discharging from full to empty battery. An unrecoverable anomaly did also conclude the experiment. These anomalies mostly involved connectivity issues at the beginning of the experiment, shortly after the device was moved in the cold chamber. On a few occasions – mostly on the older models (i.e. S2, Nexus and S3) – the logger was terminated by the operating system due to low system resources. This only occurred when the CPU usage was set to high (100%), meaning that the operating system likely needed to reallocate system resources for a higher priority system process.

Each device was monitored from the adjacent control room using the live information feed of the local server. From this information feed, both anomalies and successful exit conditions (e.g. shutdown) were observed and acted upon. When the experiment was concluded, the device was moved to the control room and plugged to a charger. Once charged to a battery capacity at which the device was expected to stay on (around 10%–20%), the charger was unplugged from the device's USB port, which was then used to transfer the logfile from internal storage to the server. After the logfile was stored alongside its remote equivalent, a few in-built applications (like the camera) were tested on the device to confirm it was intact and functioning properly. Finally, the device was plugged back to the charger and left to warm up for another round.

Measured CPU load. Although the experiment devices were factory reset and any unnecessary software was removed from the devices, some background CPU load still remains for each device as part of system services or bloatware applications. After conducting the experiment, the CPU loads were observed and the artificially generated 50% CPU load was observed to behave similar to the “CPU idle” state. In these two conditions the average CPU load is 21.4% (SD = 16.7%, IQR = 8.9–29.2). With the “CPU 100%” condition the CPU load remained within 95% and 100% load. As it was difficult to accurately differentiate between CPU idle and CPU 50% conditions, it was decided to categorise the CPU load to two conditions: *idle* (below 95%) and *high* (from 95% to 100%).

Table 5
Collected data variables.

Variable	Format	Source
Ambient temp	C°	Room
CPU usage	0–100%	Device
Battery level	0–100%	Device
Battery temp	C°	Device
Battery voltage	mV	Device
Battery current	mA	Device
Battery health	nominal	Device

Table 6
Regression function validation results, scores indicate M \pm SD of the absolute error, in degrees Celsius.

	0C	–10C	–20C	–30C
Linear model	5.32 \pm 3.66	6.28 \pm 4.9	8.15 \pm 5.32	6.48 \pm 3.57
Robust regression (rlm)	5.31 \pm 3.69	6.21 \pm 5.0	8.14 \pm 5.34	6.48 \pm 3.57
Local polynomial regression (loess)	4.50 \pm 3.72	5.18 \pm 4.0	6.07 \pm 3.36	5.09 \pm 2.50
Generalised additive model (gam)	5.32 \pm 3.67	6.29 \pm 4.9	8.15 \pm 5.32	6.49 \pm 3.57

3.3. Collected data

The study consists of 63 full discharging cycles from full battery until shutdown and 18 interrupted test cycles (reasons for termination discusses in Section 3.2). Experiments were running within two weeks in March 2017. Table 5 presents the variables that were collected. These include CPU usage, battery level, battery’s inner temperature in Celsius, battery voltage (mV) and battery current (mA), all read from the Android and iOS APIs. In addition, we had the ambient temperature collected by the room control (see Section 3.1) and a nominal variable called ‘battery health’, the software-based estimation of battery’s state provided by the device API. This “battery health” includes categorical labels such as ‘good’, ‘overheated’, ‘cold’, and ‘error’, which we further use as a baseline for our analysis.

For majority of the measurement cases, we analyse the first 5000 s as the internal battery behaviour (voltage and temperature) tends to plateau within this time frame. Additionally, it was observed that this duration is large enough to identify major differences in battery behaviour between the conditions. The battery temperature typically plateaus after 1000 s (roughly 15 min) of decline. The effect of thermal energy caused by the high CPU load greatly influences the level at which the battery temperature levels out (as seen in Fig. 3(a) through Fig. 3(d)). As the room temperature lowers, the influence of the heat caused by the CPU also diminishes, resulting in a significantly lower plateau temperature level in comparison to the ambient temperature (the lower the ambient temperature, the lower the plateau level).

4. Experimental results


4.1. Battery level

First, we investigate what influence the change in room temperature has on the on-device battery levels. Using a factorial ANOVA we can observe impact of room temperature, CPU usage, and voltage on the battery level. Lower room temperature drains the battery faster ($p < .005$, $F = 6.38$) and causes changes in the battery voltage ($p < .005$, $F = 250.79$), which likely is the main culprit behind higher battery drainage. Fig. 4 showcases the differences in battery level at different condition. The quick decline in battery level is most evident at -30 °C, which was the most extreme temperature. The CPU load also impact the battery decline rate ($p < .05$, $F = 4.03$).


We further tested the influence of the surrounding temperature on voltage alone with a one-way ANOVA and revealed a stronger relationship between the two ($p < .005$, $F = 16.23$). Lastly, using ANOVA we observe that all three (room temperature, battery temperature, and CPU load) have a significant ($p < .005$) impact on battery level decline. Similar behaviour exists for the tested iPhone model, as both the ambient temperature (and associated battery temperature), and voltage impact how quickly the battery discharges ($p < .005$).

4.1.1. Validating regression method


For the visualisation method in Fig. 3, Fig. 4, and for the interpolation method used later for modelling battery health in chapter 4.5, we evaluated the use of different regression methods. Namely standard linear regression model (lm), robust regression (rlm), local polynomial regression (loess), and generalised additive model (gam). We evaluated the different methods based on propagation of absolute error over the dataset. The collected results indicate the loess method is the best choice for both visualisation and modelling purposes, with lowest absolute error in all four temperatures. The complete results are shown in Table 6.




(a) Changes in battery temperature in a room temperature of 0C



(b) Changes in battery temperature in a room temperature of -10C



(c) Changes in battery temperature in a room temperature of -20C



(d) Changes in battery temperature in a room temperature of -30C

Fig. 3. Aggregated changes in battery temperature according to the four distinct ambient temperatures used in the experiments, using the local polynomial regression (loess) regression function [29]. Lines denote means of the regression function and bands denote 95% confidence interval.

4.2. Battery malfunctions

The reported preferred operating temperature for an Android device is from 0 °C to +35 °C, but they are seen to be able to operate in colder temperatures. None of the experimental Android devices malfunctioned in temperatures above -20 °C. Android OS includes a self-contained battery health monitoring software, which is then customised per each vendor. The health categories and the scale used to report different values of battery health are determined by the software, and thus, the vendor of each device. Using this battery health functionality, only two of the older devices

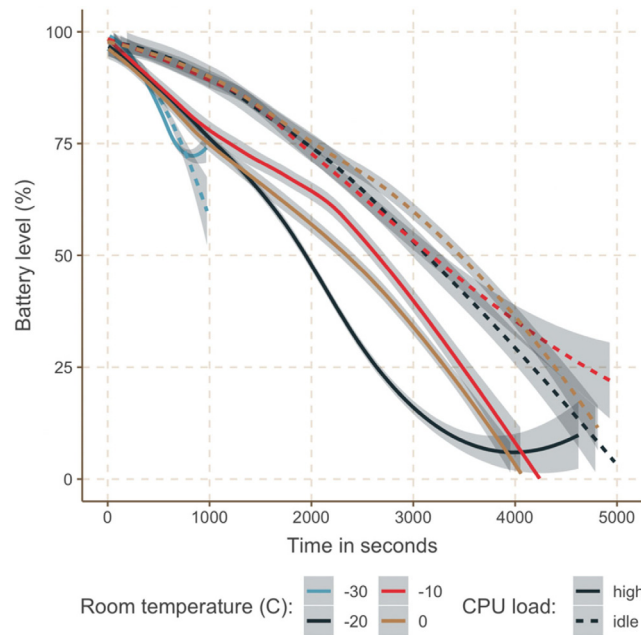


Fig. 4. Differences in battery level decline in different room temperatures according to local polynomial regression (loess [29]). Lines denote means and bands denote 95% confidence interval. As sample sizes are reduced towards end of each test, and tests end at different times (due to device malfunction), variance increased towards the end of each test.

(Samsung Galaxy I9100 and Nexus I9250) reported ‘cold’ battery health, typically once reaching battery temperatures close to -5 to -10 °C, but also in the ambient temperature of -10 °C with idle CPU load. Other devices reported the default value ‘good’ throughout the experiments although the battery temperature dropped similarly across the devices, and all devices suffered similarly from the most extreme temperature, as shown later. This behaviour highlights the problem with smartphones not being aware of external conditions that could negatively influence the device battery health and lifetime, or at extreme cases, the device’s reliability.

When exposed to temperature of -30 °C, the temperature lowered at an accelerated pace compared to higher temperatures as expected. More importantly, all tested devices automatically shut down (i.e. crashed due to a malfunction) before reaching zero battery level. Typical crash time was 575 s with median 514 s, standard deviation 238 s, and inner quartile range 473–714 s, as shown in Fig. 8(b). The devices shut down quicker when the CPU was idle (569.11s compared to 589.33 s), but not by a statistically significant margin ($p = .92$). Even though the CPU generates heat and there is a distinction in measured device temperature between idle and high CPU load, it is not enough to heat the device battery sufficiently to prevent crashes.

Again, for the tested iPhone model the crashes occurred similarly at -30 °C, after reaching a battery temperature of around -15 °C. The crashes occurred more rapidly than with the Android models ($M = 488$, $p < .05$). Alarmingly the iPhone also shut down at -20 °C ambient temperature, where its internal battery temperature barely touched -10 °C. This type of behaviour has been reported repeatedly in media and consumer feedback, and the more recent iPhone models contain a warning system for extreme (above $+35$ °C or below 0 °C) temperatures. Without appropriate warning systems, end-user devices run a risk of long-term damage due to, e.g., liquid condensation or other problems related to temperature shock.

4.3. Modelling battery health awareness in extreme conditions

As our study revealed, nowhere close all the smartphone models can appropriately showcase battery health and conditions in extreme temperatures. In our experiments, only two smartphones – Samsung I9100 (S2) and I9250 (Nexus) – did report ‘cold’ battery health outside their operating temperatures. Distribution of ‘cold’ and ‘good’ battery reports of these models is shown in Fig. 8(a). Next, we present a predictive model (build with data from I9100 and I9250) that is then be applied as a generic model to all the other devices.

Preprocessing for modelling. As discussed in Section 3.3, some on-device measurements (battery level, temperature, and voltage) are reported once every 60 s. For modelling purposes, the delay between these measurements is too high, thus we first interpolate the missing data. We apply the R smooth.spline function to each of the three variables aligned with the stored timestamp. We use a smoothing parameter of 0.9 (in a 0..1 range) to ensure sufficient curvature and distinction from the step-like form of the raw dataset.

Feature extraction. As our hypothesis on the nature of the battery health reporting system relies on current values and the aggregated change over time to better understand speed of change and the short-term historical behaviour of the battery temperature, we use a following set of features for modelling, with the number of extracted features in parentheses:

1. Current CPU load (1), change in CPU load over a time window of 1, 5, 10, 30, and 60 s (5)
2. Current battery level (1) and change in previous time windows (5)
3. Current battery voltage (1) and change in previous time windows (5)

This brings the total number of model features to 18. We use the reported battery health as the label variable. The labels consist of 46,687 entries of 'good', 21,570 entries of 'cold', and 13 entries of 'error', with a total training data size of 68,2 K and the training data being partly biased to 'good' (68.39% of samples). We use a randomised 10-fold cross validation for our evaluation. This can lead to biased training data sets in some of the folds but is generally seen as robust method for evaluation typical classifiers.

Results. As seen in Fig. 8(a), both devices seem to implement rather simplified thresholds for reporting 'cold' battery health, although I9100 offers both 'good' and 'cold' in the 0...–10 °C range. The inner hardware layout or component type also seem to significantly affect how low and how rapidly the battery temperature declines. However, a simple threshold would obviously not function as a general model, thus we test three different classifiers and evaluate their performance using 10-fold cross validation. Initial Naive Bayes function proves ineffective with 58,25% performance, but when exploring both logistic regression model and the Random Forest classifier, the classifier performance rapidly increases to 99,96% (Logistic) and 99,99% (Random Forest). While such high accuracy tends to raise eyebrows, we must realise that vast majority of the training cases belong rather explicitly to either 'cold' or 'good' class. The edge cases for our generic model are then effectively solved by the combination of decline rates in battery temperature, voltage, and level.

Using this logistic regression model on the rest of the dataset – the set of devices that solely report 'good' battery health – the model implicates that in 27.6% of the cases the device should, in fact, report the battery health 'cold' (false negative or a Type II error). No Type I errors were reported in the model. Thus, for majority of the models used in this experiment, the on-board battery health monitoring software is critically lacking.

4.4. Modelling for crash prevention

Similarly, to modelling the battery health status, we also aimed to go a step further and analyse the feasibility of predicting device crashes when exposed to harsh cold temperatures. As mentioned, all device models eventually crashed in the coldest (–30 °C) experimental condition (see Fig. 8(b)), thus we are able to use all 427 K data points for training and testing the model with 10-fold cross validation.

To measure the *alert level for an upcoming crash event*, we use a 5-tier system similar to ones used in e.g., US homeland security. The tiers are 'low', 'guarded', 'elevated', 'high', and 'severe'. A threshold for 'low' is no upcoming crash event, and a potential crash even within 10 min, 5 min, 120 s, and 60 s, respectively. Considering the full dataset 97.9% (418 K) entries belong in the 'low' crash level, thus when evaluating the model this should be taken into consideration. If the resulting model would be clearly overfitted to a particular class, we could down-sample and balance the dataset – but this ends up not being needed.

Modelling concept. The key concept of the modelling approach is to create a warning system that can detect future crash events. We use the 5-tier alert level system as the output of the model, thus the model could provide the user with an expected usage time according to the 5-tier nominal output variable (e.g., 120 s if the model outputs a 'high' alert level) before the device crashes if no steps are taken to protect the device. This would provide the user with time to plan his actions to ensure device operation in cold or extreme conditions. The model's features (as described in paragraph *Feature extraction*) can be continuously monitored and generated using on-device sensors – namely the battery sensor and the BatteryManager.¹

Model evaluation. Again, we initially evaluate the model using Naive Bayes as a baseline (with 68.31% accuracy), and then proceed to build models with logistic regression (98.48% accuracy) and Random Forest (99.97%), evaluated using a standard 10-fold cross validation and default parameters from the Weka package.² The confusion matrix from the Random Forest is presented in Table 7. When considering the 'high' and 'severe' classes, the matrix showcases high recall (936/949 and 1016/1028, respectively) and high precision (936/962 and 1016/1028) in both cases. A total of 26 (of 1952) Type I errors (model falsely reporting crash threat level as lower as intended and total of 12 (of 1952) Type II errors (model falsely reporting threat level as higher) were found in the model's results in these two categories.

Further evaluating the efficacy of the Random Forest model, the out-of-bag (OOB) estimate is 99.97% with a Kappa statistic of .994. Kappa statistic is a metric that compares Observed Accuracy with Expected Accuracy (random chance based on the properties of the tested dataset). The model's feature importance is measured using average impurity decrease which evaluates the effect of removing a feature on the overall performance of the model. Top-three predictor features are the current readings (battery *temperature*, battery *level*, and *voltage*, respectively), and from the time window-based features changes in battery *temperature* has the most importance, starting from the longest time window of 60 s

¹ <https://developer.android.com/training/monitoring-device-state/battery-monitoring>

² <https://www.cs.waikato.ac.nz/ml/weka/>

Table 7

Confusion matrix for a Random Forest model predicting the 'crash threat level'. Time (T) within parenthesis indicates that a crash is likely to occur within T.

Classified as:	low (N = 418K)	guarded (3.9K)	elevated (2.9K)	high (962)	severe (1028)
low (no crash threat)	418356	13	0	0	0
guarded (10 min)	22	3946	11	0	0
elevated (5 min)	2	10	2921	7	0
high (120 s)	1	0	13	936	12
severe (60 s)	0	0	3	6	1016

with highest importance (rank #4), with next importance ranks #5-#8 progressing in time order of 30 s, 10 s, 5s for temperature. The average impurity decrease metric itself does not relate to the original dataset, but the features are generally quite well balanced with the *most important* feature having a decrease of .51 (with 4004 tree nodes using that feature) to *least important* feature with a decrease of .32 (981 nodes). It is not uncommon to have near-zero decrease metric scores in some models, thus even the least important features in our model provide significant information.

Overall, we can conclude that the model would be effective in differentiating the situations where the device is under a threat of crashing from normal usage.

Extending Battery Health Modelling & in-the-wild. These simple models showcase that these features would be quite effortless to implement for current smartphones, using the on-board sensors and a rather generalised model to start with. The lack of such systems currently implemented appropriately, and the potential problems and harm battery malfunction can cause we wish to bring this issue to the attention of the research community as well as the end-users and mobile manufacturers. However, although the occurrence of a device crash due to extreme conditions is an important event to avoid, and in some cases can even be life threatening, it is likely an event which happens very infrequently and only in very specific conditions. Investigating these occurrences *in-the-wild* in a scope where these events could be reliably observed (and potentially avoided) would require vast research projects where such battery-monitoring applications would be propagated to a very wide audience. Alternatively, we could purposely expose devices or personnel to these conditions *when such conditions occur naturally*, but even these approaches would not be viable without long-term research ventures (that would likely span over several winters).

4.5. Lasting effects on battery health

Lastly, we explored the notion that a device shut down due to exposure to extreme cold temperature could cause long-term damage to the battery that could result in future malfunctions. The work of Waldmann [30] explores the effect of temperature-based ageing on Lithium-ion batteries in post-mortem analysis, i.e., depleting and essentially destroying the battery in extreme conditions and investigating how it was affected. Their results showcase that lithium batteries age and cause capacity decline significantly faster in -20C or lower temperatures due to a process known as lithium plating. Lin et al. [31] have also highlighted that subzero (C) temperatures can cause performance loss and temperatures of below 20C cause permanent capacity loss. Other battery performance degrading effects are also exhibited when lithium-based devices are used outside of their suggested operating temperature of around 25C . While our study did not analyse battery capacity, we noticed peculiar behaviour in testing for three of the devices, the I9506, Nexus 5, and I9250, at temperatures of -30C and -20C . *Note:* All three devices were tested in 0C and -10C before the set of tests presented below and behaved normally in these conditions.

Fig. 5 shows the voltage and battery temperature readings from I9506 in subsequent test scenarios. The device was always returned to room temperature and charged to full in between tests. In test #51 the device lasts until 1049s before crashing, but on the very next test #52 the device crashes after 330 s. There was a waiting time of 33 min 20 s between tests #51 and #52. Tests #53 and #54 were conducted later on the same day, with a similar 30-minute pause in between. Test #53 mimics the behaviour of #52 with a rapid failure after a sharp decline in battery voltage, but on test #54 the voltage decline is slower, and the device again lasts over the 1000-second mark before crashing. After these set of tests test #56 was conducted on the same device and the battery voltage seemed to behaviour erratically, with sharp dips after the battery core temperature had reached -20C .

While the I9506 did not clearly describe any short-term harm for the battery, Fig. 6 shows how subsequent testing of the Nexus 5 in just -20C environment lead to declining reliability and caused the Nexus 5 to crash even in -10C temperature when under heavy load (*high CPU load*). The cause for the crashes was not the temperature changes but rather the battery voltage which seemed to have be due to initial test #74. In subsequent tests #75, #76, and #77 the battery voltage dips very rapidly and causes a quick crash. Similarly, in test #79 *high CPU load* (faster battery discharge) causes decline in battery voltage and causes the device to crash even though the temperature of the battery stays well above 25C (see Fig. 7).

Lastly, the I9250 crashed on all three tests at -30C , lasting for 792 s on the first test #25, but only for 496 s and 535 s on the following two tests, #26 and #27. It should also be noted that the battery temperature did not reach far below 0C but rather battery voltage seemed to be the culprit behind the crashes. The I9250 did not crash in -20C temperatures,

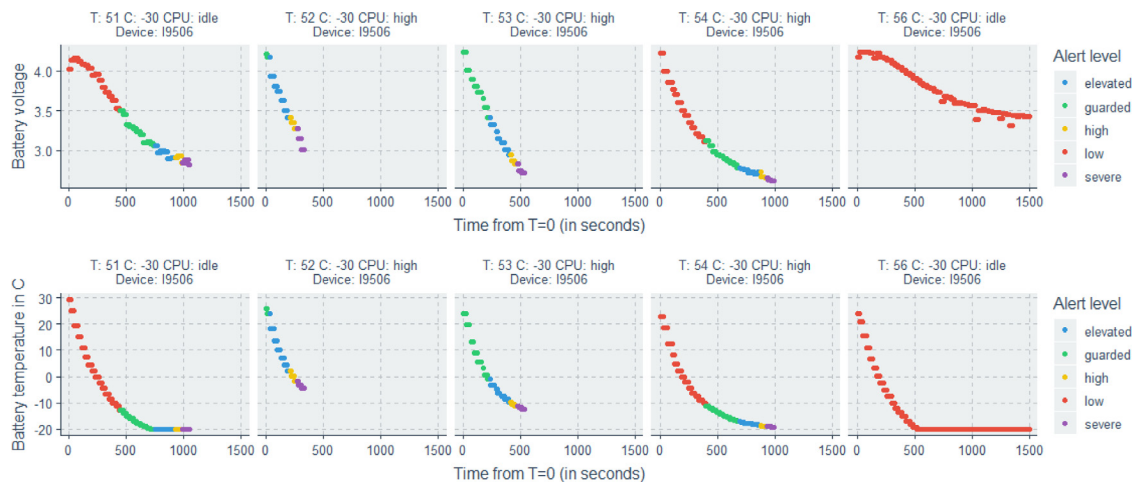


Fig. 5. Samsung i9506 crash times and ground truth alert levels in selected test cases (labelled 'T' in figure headings).




Fig. 6. Nexus 5 crash times and ground truth alert levels in selected test cases (labelled 'T' in figure headings).

but similar to the I9506 the battery voltage started to act erratically, as shown in tests #29, #30, and #31. This behaviour could be due to damage caused by the extreme temperatures of tests #25–#27.

Using factorial ANOVA we found no significant evidence ($p = .06$, $F = .71$) that this is case nor that the device model influences ($p = .12$, $F = .36$) how quickly the device shuts down. All devices brought to the coldest temperature crashed on at least one occasion.

Due to the constraints of the facilities used, this type of testing would have to be done in a less controlled environment, and these experiments would be subject to, e.g., normal changes in outside temperature. Experimenting with long-term wear, especially from the perspective of temperature shock-related wear, would be important for future research, as battery wear can be associated with battery malfunctions (e.g., short circuits) that can cause direct harm to end-users [32].

5. Discussion

It can be widely questioned why anyone would use their smartphones in extreme subzero temperatures. However, our current highly smartphone-dependent lifestyles and, to say the least, any sudden emergency situations can require battery life and some level of performance being left in the device even after shortly visiting, working, or travelling outside in cold. Based on our empirical measurements in this study, it is clear that $-30\text{ }^{\circ}\text{C}$ is too much for a smartphone and will irrevocably and rapidly kill the device (within 5–10 min at worst). However, $-30\text{ }^{\circ}\text{C}$ or colder is definitely usual in the Northern parts of Scandinavia, Russia, and Canada during the winter days. Even in less extreme conditions, the effect of cold ambient temperature causes the battery level to decline in an accelerated fashion.

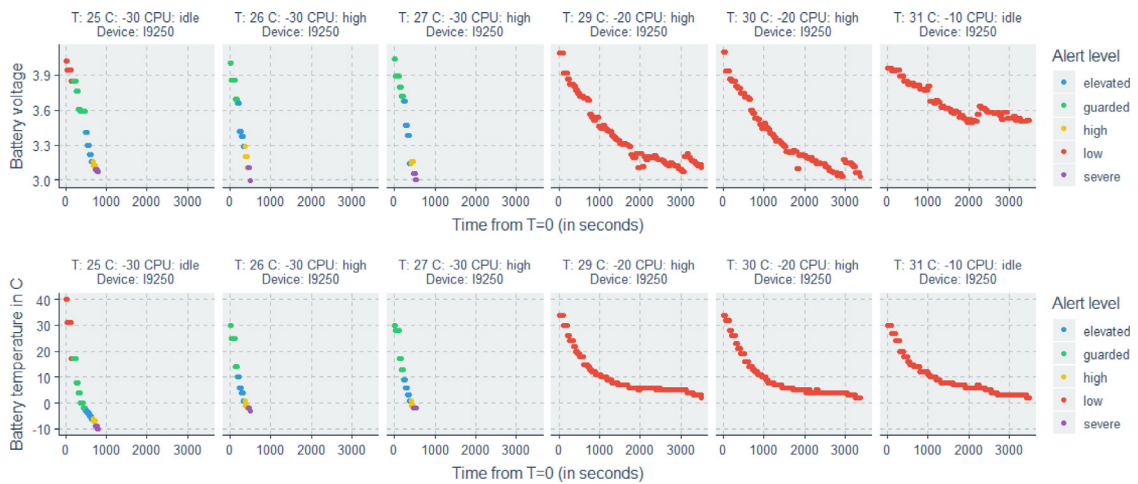
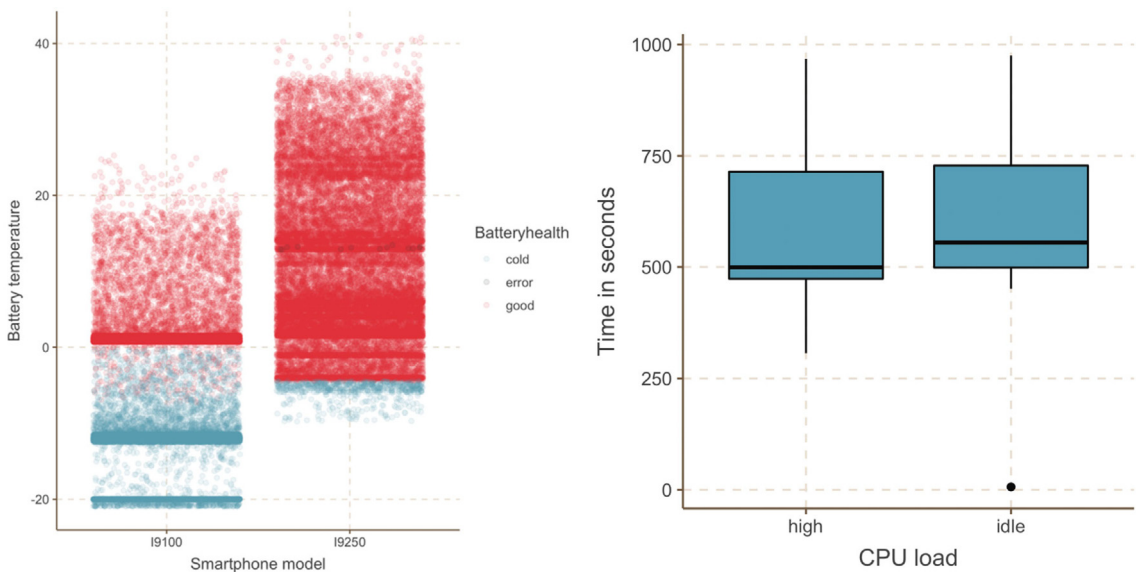


Fig. 7. Samsung i9250 crash times and ground truth alert levels in selected test cases (labelled 'T' in figure headings).



(a) 'Good' and 'cold' battery health.

(b) Crash times in -30°C .

Fig. 8. (a) Battery health reports from the two devices (Galaxy I9100 and Nexus I9250) that make a distinction between 'good' and 'cold' battery health. (b) Crash times in -30°C under 'high' CPU load and 'idle'.

In somehow milder but still definitely cold temperatures, particularly in the -20°C region, high CPU use causes faster battery discharge similar to what happens in "normal" operating temperatures. However, the CPU generating internal heat to warm up otherwise cooling battery could potentially be used to resist effects of subzero ambient temperature. Thus, it can be speculated if CPU management could ensure longer device lifetime. However, such an engineering should require careful balancing between the inside temperature compensation and faster energy use caused by higher CPU load.

The impact of different ambient temperatures on smartphones battery temperature is also evident in our experiment. Discrepancy between ambient and battery temperature behaves very differently as the ambient temperature declines. This could have significant impacts on any crowdsensing solutions aimed at measuring e.g., urban air temperatures [11], as the smartphone-based temperature readings can become unreliable outside typical operating temperatures. The models created in [11,12] would likely require further datasets in cold climates and sub-zero temperatures to extend their models to fit such environments.

Different potential solutions exist for protecting the smartphone from such ambient effects of outdoors in subzero temperatures. Extensive carriage solutions (such as keeping the phone close to the body or inside pocket of the winter clothes) clearly limit the actual usage of the smartphone, even if it can help lasting the battery and keeping the device warmer outdoors. Most of the smartphone cover cases are designed for protecting the device from physical shocks, scratches, or humidity. Their thermal insulation should be more widely tested and analysed to see if they have any help against subzero temperatures. However, it is clear that the cheapest plastic cover cases do not add enough if the device's own protective shields do not do that already. The device could also attempt to protect itself by appropriately monitoring its surroundings, and our findings reveal that with most device models these capabilities are severely lacking. Several battery management applications exist that come either pre-installed or can be installed through application repositories, but these rarely take into account temperature as the battery condition. As shown in our data, some vendor-provided solutions – specifically the Android OS battery health monitoring – is critically flawed. This shows an important future research goal of extending the set of devices tested. One solution can also simply be a third-party application that uses models similar to ones proposed in this paper. These applications could prevent the battery from reaching temperatures that can lead to long-term damage and will help sustain the device's reliability, in addition to alerting the user in extreme conditions.

Limitations. As previous studies in the field have been mostly oriented towards human–computer interaction, we cannot directly compare our results and methodology against any known work, but instead try to establish a baseline for temperature measurements on smartphones and unveil the implications of cold temperature on their behaviour. Because our study involved extremely cold temperatures, the target devices were at a constant risk of hardware breakdown. This imposed a financial constraint which limited the selection to slightly older devices. In order to maintain a constant temperature and climate conditions, the study was conducted in a climatic cold chamber. Although this allowed us to control the environmental setting, we acknowledge that the simulation may not fully reflect natural outdoor conditions.

Future work. This paper presents initial results in how an on-board smartphone (opposed to detached factory test) battery behaves in subzero and extremely cold temperatures. We used relatively old models in our experiment, which all used Lithium-Ion batteries. Newer high-end smartphone models typically use Lithium-Polymer. While the functionality of these battery types is generally similar [33], it is clear that the behaviour of these polymer-based batteries in extreme temperatures should be further investigated. The results reveal several insights which would benefit from further research, including a longitudinal analysis and real-world data from actual users using their smartphones in cold environments. Crowdsensing studies have demonstrated to be an effective tool for collecting large-scale real-world data [3], e.g., temperature. Our laboratory-driven study has provided a controlled baseline that could further used to validate crowdsensed data. In addition, the battery level seems to decline rapidly when there is a significant change in the on-board battery temperature. Whether this behaviour occurs when temperature rises, is unclear, however. Lastly, the crash behaviour could be studied further by, e.g., whether returning the device to a warmer environment prior to typical crash time would prevent crashes. Investigating the impact of transitions from indoors to outdoors would be an interesting research topic. In addition, with a well-defined research environment and experimental setup, it becomes interesting to study also other ends of the extreme temperatures: extreme heat and humidity. s

6. Conclusions

This paper presents an overview of smartphones being exposed to subzero and extreme cold temperatures. We used a medical climate chamber to simulate cold climate conditions in a clinical laboratory setting. This allowed us to create fair conditions for running the experiment by maintaining constant temperature and controlling climate factors such as precipitations, wind chill, wind speed, and humidity.

The results showcase how rapidly the on-board battery temperature diminishes, and how higher CPU loads can withhold higher device temperatures during use. Even brief exposures to -30°C (or lower) can rapidly – even in less than 15 min – cause a typical Android or iOS device to shut down or even freeze due to extreme conditions. This behaviour or the exposure to cold temperatures is not reported by newer smartphone models, even 27% of potential exposures left unobserved by the operating system. However, we found out that predicting battery health awareness and even operating system's crashes is definitely effortless task with simple regression models gaining high accuracy. This begs to question the usability of the embedded battery health sensor, and missing utilisation of battery temperature sensor to provide warnings in extremely cold environments.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

- [1] J. Jaguemont, L. Boulon, Y. Dubé, A comprehensive review of lithium-ion batteries used in hybrid and electric vehicles at cold temperatures, *Appl. Energy* 164 (2016) 99–114.
- [2] Z. Sarsenbayeva, N. Van Berkel, A. Visuri, S. Rissanen, H. Rintamäki, V. Kostakos, J. Goncalves, Sensing cold-induced situational impairments in mobile interaction using battery temperature, *ACM Interact. Mob., Wearable Ubiquitous Technol.* 1 (3) (2017).
- [3] E. Peltonen, E. Lagerspetz, P. Nurmi, S. Tarkoma, Energy modeling of system settings: A crowdsourced approach, in: *IEEE Int. Conf. on Pervasive Computing and Communications*, 2015, pp. 37–45.
- [4] A.J. McMichael, R.E. Woodruff, S. Hales, Climate change and human health: present and future risks, *Lancet* 367 (9513) (2006).
- [5] A. Vlahinos, A.A. Pesaran, Energy efficient battery heating in cold climates, in: *SAE Technical Paper, Tech. Rep.*, 2002.
- [6] T.M. Mäkinen, L.A. Palinkas, D.L. Reeves, T. Pääkkönen, H. Rintamäki, J. Leppäluoto, J. Hassi, Effect of repeated exposures to cold on cognitive performance in humans, *Physiol. Behav.* 87 (1) (2006).
- [7] F. Mueller, S.J. Pell, Technology meets adventure: learnings from an earthquake-interrupted Mt. Everest expedition, in: *ACM Int. Joint Conf. on Pervasive and Ubiquitous Computing*, ACM, 2016, pp. 817–828.
- [8] H. Virokannas, Thermal responses to light, moderate and heavy daily outdoor work in cold weather, *Eur. J. Appl. Physiol.* 72 (5–6) (1996) 483–489.
- [9] Z. Sarsenbayeva, J. Goncalves, J. García, S. Klakegg, S. Rissanen, H. Rintamäki, J. Hannu, V. Kostakos, Situational impairments to mobile interaction in cold environments, in: *ACM Int. Joint Conf. on Pervasive and Ubiquitous Computing*, ACM, 2016, pp. 85–96.
- [10] Y. Liu, J. Goncalves, D. Ferreira, S. Hosio, V. Kostakos, Identity crisis of ubicomp?: mapping 15 years of the field's development and paradigm change, in: *ACM Int. Joint Conf. on Pervasive and Ubiquitous Computing*, ACM, 2014, pp. 75–86.
- [11] A. Overeem, J.R. Robinson, H. Leijnse, G.-J. Steeneveld, B.P. Horn, R. Uijlenhoet, Crowdsourcing urban air temperatures from smartphone battery temperatures, *Geophys. Res. Lett.* 40 (15) (2013) 4081–4085.
- [12] A. Droste, J.Pape, A. Overeem, H. Leijnse, G. Steeneveld, A. Delden, R. Uijlenhoet, Crowdsourcing urban air temperatures through smartphone battery temperatures in São Paulo, Brazil, *J. Atmos. Ocean. Technol.* 34 (9) (2017) 1853–1866.
- [13] Q. Xie, J. Kim, Y. Wang, D. Shin, N. Chang, M. Pedram, Dynamic thermal management in mobile devices considering the thermal coupling between battery and application processor, in: *IEEE/ACM Int. Conf. on Computer-Aided Design*, 2013, pp. 242–247.
- [14] S. Gurrum, D. Edwards, T. Marchand-Golder, J. Akiyama, S. Yokoya, J. Drouard, F. Dahan, Generic thermal analysis for phone and tablet systems, in: *IEEE Electronic Components and Technology Conf.*, 2012.
- [15] Q. Xie, M.J. Dousti, M. Pedram, Terminator: A thermal simulator for smartphones producing accurate chip and skin temperature maps, in: *IEEE/ACM Int. Symp. on Low Power Electronics and Design*, 2014, pp. 117–122.
- [16] Z. Qian, Y. Li, Z. Rao, Thermal performance of lithium-ion battery thermal management system by using mini-channel cooling, *Energy Convers. Manage.* 126 (2016) 622–631.
- [17] X. Duan, G. Naterer, Heat transfer in phase change materials for thermal management of electric vehicle battery modules, *Int. J. Heat Mass Transfer* 53 (23–24) (2010) 5176–5182.
- [18] C.-Y. Wang, G. Zhang, S. Ge, T. Xu, Y. Ji, X.-G. Yang, Y. Leng, Lithium-ion battery structure that self-heats at low temperatures, *Nature* 529 (7587) (2016) 515.
- [19] S. Aram, A. Troiano, E. Pasero, Environment sensing using smartphone, in: *IEEE Sensors Applications Symposium*, 2012, pp. 1–4.
- [20] J. Jaguemont, L. Boulon, Y. Dubé, Characterization and modeling of a hybrid-electric-vehicle lithium-ion battery pack at low temperatures, *IEEE Trans. Veh. Technol.* 65 (1) (2015) 1–14.
- [21] L.A. Palinkas, T.M. Mäkinen, T. Pääkkönen, H. Rintamäki, J. Leppäluoto, J. Hassi, Influence of seasonally adjusted exposure to cold and darkness on cognitive performance in circumpolar residents, *Scand. J. Psychol.* 46 (3) (2005) 239–246.
- [22] M.B. Kjærgaard, H. Blunck, Unsupervised power profiling for mobile devices, in: *Int. Conf. on Mobile and Ubiquitous Systems: Computing, Networking, and Services*, Springer, 2011, pp. 138–149.
- [23] F. Xu, Y. Liu, Q. Li, Y. Zhang, Fast self-constructive power modeling of smartphones based on battery voltage dynamics, in: *USENIX Symp. on Net. Systems Design and Impl.*, 2013, pp. 43–55.
- [24] A. Pathak, Y.C. Hu, M. Zhang, Where is the energy spent inside my app?: fine grained energy accounting on smartphones with eprof, in: *ACM European Conf. on Computer Systems*, 2012, pp. 29–42.
- [25] L. Zhang, B. Tiwana, Z. Qian, Z. Wang, R.P. Dick, Z.M. Mao, L. Yang, Accurate online power estimation and automatic battery behavior based power model generation for smartphones, in: *IEEE/ACM/IFIP Int. Conf. on Hardware/software codesign and system synthesis*, 2010.
- [26] Y. Liu, C. Xu, S.-C. Cheung, Where has my battery gone? Finding sensor related energy black holes in smartphone applications, in: *IEEE Int. Conf. on Pervasive Computing and Communications*, 2013, pp. 2–10.
- [27] X. Ma, P. Huang, X. Jin, P. Wang, S. Park, D. Shen, Y. Zhou, L.K. Saul, G.M. Voelker, eDoctor: Automatically diagnosing abnormal battery drain issues on smartphones, in: *USENIX Symp. on Net. Systems Design and Impl.*, 2013, pp. 57–70.
- [28] A.J. Oliner, A.P. Iyer, I. Stoica, E. Lagerspetz, S. Tarkoma, Carat: Collaborative energy diagnosis for mobile devices, in: *ACM Conference on Embedded Networked Sensor Systems*, ACM, 2013, p. 10.
- [29] R.A. Cohen, An introduction to proc loess for local regression, in: *Proceedings of the Twenty-Fourth Annual SAS Users Group International Conference, Paper, Vol. 273*, Citeseer, 1999.
- [30] T. Waldmann, M. Wilka, M. Kasper, M. Fleischhammer, M. Wohlfahrt-Mehrens, Temperature dependent ageing mechanisms in lithium-ion batteries—a post-mortem study, *J. Power Sources* 262 (2014) 129–135.
- [31] H.-P. Lin, D. Chua, M. Salomon, H. Shiao, M. Hendrickson, E. Plichta, S. Slane, Low-temperature behavior of li-ion cells, *Electrochem. Solid State Lett.* 4 (6) (2001) A71.
- [32] P.J. Mankowski, J. Kanevsky, P. Bakirtzian, S. Cugno, Cellular phone collateral damage: A review of burns associated with lithium battery powered mobile devices, *Burns* 42 (4) (2016).
- [33] Y. Miao, P. Hynan, A. Von Jouanne, A. Yokochi, Current li-ion battery technologies in electric vehicles and opportunities for advancements, *Energies* 12 (6) (2019) 1074.