



Adaptive Context-Aware Energy Optimization for Services on Mobile Devices with Use of Machine Learning

Piotr Nawrocki¹ · Bartłomiej Sniezynski¹

Published online: 13 August 2020
© The Author(s) 2020

Abstract

In this paper we present an original adaptive task scheduling system, which optimizes the energy consumption of mobile devices using machine learning mechanisms and context information. The system learns how to allocate resources appropriately: how to schedule services/tasks optimally between the device and the cloud, which is especially important in mobile systems. Decisions are made taking the context into account (e.g. network connection type, location, potential time and cost of executing the application or service). In this study, a supervised learning agent architecture and service selection algorithm are proposed to solve this problem. Adaptation is performed online, on a mobile device. Information about the context, task description, the decision made and its results such as power consumption are stored and constitute training data for a supervised learning algorithm, which updates the knowledge used to determine the optimal location for the execution of a given type of task. To verify the solution proposed, appropriate software has been developed and a series of experiments have been conducted. Results show that as a result of the experience gathered and the learning process performed, the decision module has become more efficient in assigning the task to either the mobile device or cloud resources.

Keywords Adaptation · Context-aware system · Energy optimization · Machine learning · Mobile cloud computing

The research presented in this paper received financial support from AGH University of Science and Technology Statutory Project. We would like to thank Karol Wojcik and Marcin Bogusz for their assistance with implementation and testing. Neither the entire paper nor any part of its content has been published or has been accepted for publication elsewhere. It has not been submitted to any other journal.

✉ Piotr Nawrocki
piotr.nawrocki@agh.edu.pl

Bartłomiej Sniezynski
bartlomiej.sniezynski@agh.edu.pl

¹ Department of Computer Science, Faculty of Computer Science, Electronics and Telecommunications, AGH University of Science and Technology, al. A. Mickiewicza 30, 30-059 Kraków, Poland

1 Introduction

The rapid development of mobile devices and the growing importance of applications and services that run on these devices has resulted in a need to pay more attention to the quality parameters associated with the use of such solutions. So far, in the context of the operating quality of applications and services of this type [1], primarily QoS (Quality of Service) parameters were considered that cover the network connection status (including delay, jitter and bandwidth) and QoE (Quality of Experience) parameters that show the level of user satisfaction with a given application or service. It should be noted that in the context of mobile devices using battery power, a prerequisite for maintaining the highest level of both QoS and QoE parameters is ensuring the longest possible uptime of the device. It can even be claimed that ensuring the longest possible uptime of the mobile device (alongside the applications and services running on it) from the point of view of conserving power is in fact a QoE parameter because it increases the level of user satisfaction with mobile applications and services.

Increasing mobile device uptime is possible through optimizing power consumption while preserving the best possible quality parameters of mobile services and applications. Such optimization can be implemented at the software or hardware levels and should take into account the context in which the mobile device operates, including network connection quality, location and potential time and cost of executing the application or service. The use of context information may allow the adaptation of mobile services and applications to prevailing conditions in order to improve the quality parameters (including execution time) and to optimize power consumption. The adaptation process in the context of optimizing power consumption can be implemented with respect to the mobile device itself and also to the mobile services and applications running on the device. Methods allowing for adaptations of this type often use remote resources such as cloud computing (the Mobile Cloud Computing concept [2]) or another mobile device with appropriate resources (the cloudlets concept), to which applications/services or their components are offloaded in order to optimize the operation of the mobile device [3]. The choice of when and what to offload from the mobile device can be made offline during the software development process or dynamically when the device is working. Adaptation makes it possible to reduce the time and costs of executing applications/services on mobile devices and to optimize their power consumption online.

In this paper we present our original concept for an adaptive system for optimizing the power consumption of mobile devices using context-based and machine learning mechanisms. Having analyzed existing solutions, we found out that there are no systems that would enable the dynamic, online adaptation of mobile applications/services by using machine learning algorithms while simultaneously accounting for the context in which the mobile device is located. A few articles describe solutions that use machine learning algorithms but these have some limitations such as requiring the use of models developed previously in offline mode or not taking into account specific applications/services when optimizing the power consumption of a mobile device. Our innovative solution works fully online on the mobile device being optimized and enables dynamic adaptation using the context in which applications/services are executed on the mobile device. It enables the costs and time required to execute mobile applications/services to be reduced and helps to optimize the power consumption of the mobile device.

The classic approach to online learning is based on reinforcement learning [4]. Our solution is based on supervised learning, like in [5], in which it is executed offline. We

show that it is possible to apply it online. Instead of preparing the training data set once at the beginning, the training data can be automatically extended and the learned knowledge updated while it is in use, just like in reinforcement learning. It is a novel approach to online learning. Our results from other domains demonstrate that fewer trials are required to find a good policy than in the case of reinforcement learning, especially in a complex environment (see e.g. [6]). This research shows that this type of learning is applicable to practical problems that are encountered in mobile computing. As a result, knowledge specific to every device may be efficiently learned locally in this device even while it is in use. There is no need to prepare a separate strategy for every device type offline. This also resolves scalability and privacy issues.

The structure of the article is as follows: Sect. 2 presents the analysis of research in the field of power optimization in the context of mobile devices, Sect. 3 presents solutions in the field of machine learning on mobile devices, Sect. 4 describes the adaptive power optimization system developed for mobile devices using context information, Sect. 5 introduces the results of the experiments conducted and Sect. 6 contains conclusions.

2 Related Work

Issues related to power management are becoming increasingly relevant, *inter alia* in the context of modern distributed systems, including those using virtualization and cloud computing. In [7], the authors present an analysis of power-saving techniques and examine the capabilities of machine learning in automatic power management systems. However, the article lacks a broader discussion of machine learning algorithms and aspects of possible adaptation. The analysis only considers desktop and similar systems and does not cover mobile devices, which have become an important part of modern distributed systems.

The power consumption aspect has been very important since mobile devices, including mobile phones, first came into existence. A lot of papers have been published on this matter including [8, 9] and some recent studies where authors present a general analysis of power management in the context of mobile devices [10] and energy saving strategies in the context of mobile device applications [11]. In [9], the authors present different methods described in literature that allow for increasing the energy efficiency of mobile devices at the software and hardware levels, including power management at the level of operating system, the management of sensors and communication interfaces and the use of cloud computing. [11] presents the strategies that can be implemented by the mobile application developer: Mobile Computation Offloading, sequential programs and GUI design.

Some publications such as [12] analyze in more detail the possibilities of saving mobile device power using cloud computing. The authors present an analysis of power consumption required when offloading calculations to the cloud using the network interfaces of mobile devices. They also analyze situations where using the Mobile Cloud Computing (MCC) concept may not lead to power savings. This may be related to privacy and data security considerations that necessitate more CPU usage (e.g. in the case of data encryption processes) and thus cause increased power consumption. Ensuring the reliable execution of certain services in the absence of proper communication with the cloud may also lead to increased mobile device power consumption. However, the authors only discuss theoretical considerations related to power saving in the Mobile Cloud Computing environment. Their research does not include practical tests and

discussions of possible adaptation (e.g. with the use of machine learning algorithms) to prevailing conditions in which services are executed in the context of their execution time and power consumption.

A similar concept for offloading applications from a mobile device in order to reduce power consumption and increase efficiency is presented by the authors of [13]. However, those studies do not use the Mobile Cloud Computing concept but rather the idea of cloudlets that allows the offloading of applications/services to nearby mobile resource-rich devices. The authors present the algorithm developed, which enables the offloading of data in an optimal manner, taking into account mobile device load and the availability of cloudlets. However, those studies lack practical tests that would present how the algorithm developed affects the power consumption of mobile devices and also tests that would include the context of running applications.

Also in the article [14], the authors use cloudlets (and private / public cloud servers) for multilevel full and partial offloading strategies. However, the tests of the developed solution were carried out only by simulation without using real devices.

An important article discussing a broad range of issues related to power consumption in the case of mobile applications enabling mobile e-learning is [15]. It presents research related to adaptive mobile systems that can learn while simultaneously showing how it is possible to expand these systems to make them aware of power consumption, thus allowing for power savings. The article analyzes the power management, modeling and adaptation aspects in the context of this type of system. The authors also conducted a critical analysis of existing constraints and the possibility of accounting for energy aspects in systems of this type. However, there is no thorough analysis of the possibilities of using machine learning algorithms that use the context of learning applications running on mobile devices and of enabling power optimization for devices of this type.

The use of mobile device communication modules has a significant impact on power consumption. [16] introduces the concept of minimizing data transmission costs in the Mobile Cloud Computing environment. The authors propose a solution that uses and at the same time extends the CloneCloud environment [17], which makes it possible to analyze the code and select only the most important elements that need to be offloaded to the cloud. The tests conducted, *inter alia* with the use of face recognition applications, showed a reduction in the time required for offloading the required data to the cloud and a reduction in service execution time as well as lower power usage by mobile devices. It is an interesting solution for reducing mobile device load; however, it does not take into account such factors as the context in which the service is executed. At the same time, some elements are always offloaded to the cloud, although in some cases this may not be efficient due to the service execution time and power consumption considerations. Therefore, adding to the concept developed the possibility of using machine learning algorithms that could learn which elements of the application and when should be offloaded to the cloud would in our opinion enable the required amount of data transferred to be reduced even further and mobile device power consumption to be optimized to a greater degree.

Machine learning is a popular technique used to optimize energy consumption. It may be applied in large scale systems, e.g. for data center scheduling [5]. Berral et al. apply supervised learning to create models that predict important system parameters (power consumption levels, CPU loads etc.). The models are learned from previous system behaviors and they are used to optimize scheduling decisions. It is a similar approach to ours, but we target mobile devices and the models learned are simpler because they only predict a single value for the computational task performed. Moreover in [5], the learning process occurs offline.

Another learning strategy is used in [18]. Reinforcement learning, a typical approach to online learning, is applied for adapting the routing protocol used by an underwater sensor network with the goal being to prolong the lifetime of such networks. Simulation results show that the adaptation extends lifetime by 20 percent.

So far, few studies have been conducted that concern the analysis of power consumption and possibilities for reducing it and that would at the same time use the context in which the mobile device is operating, and machine learning algorithms that allow for power consumption optimization. One of the most important articles covering these topics is [19]. In that article, the authors analyze machine learning algorithms in the context of power saving in mobile systems. They present a concept where power saving is possible thanks to the dynamic adaptation of data transfer and network interface parameters, which is conducted automatically without user intervention. It involves the appropriate startup and shutdown of mobile device network and localization interfaces. The authors also conducted actual tests for five different user profiles and five different smartphones using the Android operating system. Each device was running the Context Logger application, which logged the context of individual mobile device users to an external server where the data were analyzed offline using an algorithm developed by the authors. Those data, combined with the context in which the device and user were operating, included details such as the day of the week, location, Wi-Fi signal strength, 3G signal strength, battery status, CPU utilization and device motion. The analysis of those data from a week-long test allowed the evaluation of user behavior, including the correlation of the need for data offloading with location. On that basis, certain user behavior patterns were determined that were associated with the context of using mobile devices and five models were proposed that defined the users' (e.g. employees' or students') behavior. At the same time, the authors analyzed five different machine learning algorithms: Linear Discriminant Analysis, Linear Logistic Regression, Non-Linear Logistic Regression with Neural Networks, K-Nearest Neighbor and Support Vector Machines, which were used in the tests. However, that analysis only concerned general use and did not take into account the character of mobile systems and solutions such as Mobile Cloud Computing. The authors also proposed a model for mobile device power consumption for devices running the Android (2.3.3) system, using the Monsoon Solutions¹ power monitor software. For this purpose, they connected the mobile device to a PC and monitored power consumption in real-time when the communication and localization interfaces were turned on (active state) and off (idle state). As a result, they determined the average power consumption for each of the interfaces. However, this method appears to have some disadvantages. The momentary radio signal strength, which is related to distance from the transmitter, directly affects the power consumption of the transmitting/receiving interfaces of the mobile device. User location and potential terrain obstacles can also influence the power consumption of the localization interface (GPS). Those aspects were not addressed by the authors at all. Power measurements and determining the model for a particular device, which is treated as a whole, exclusively under laboratory conditions does not answer the question concerning the practical characteristics of power consumption by communication interfaces of a mobile device. In this context, our research on power consumption (e.g. communication interfaces) included not only the mobile device itself but also individual services executed on it, which allowed for a more accurate analysis of energy aspects. The studies that we carried out used the PowerTutor software [20], which allows for a continuous analysis of power consumption. Using online analysis instead of predetermining a model

¹ Monsoon Solutions—www.msoon.com.

under laboratory conditions (as the authors did [19]) allows for a better inclusion of the actual context in which the mobile device is located. The authors [19] carried out a series of tests that demonstrated the effectiveness of using different machine learning algorithms to dynamically predict the energy efficiency of different device interface configurations. Among other things, they were successful in 90% of cases when predicting power consumption using support vector machines, neural networks and k-nearest neighbor algorithms. However, it appears that such research requires tests on a larger scale than involving just a few users. For such a small number of devices, the patterns discovered may not be fully useful and universal. We believe that when the operation of the entire device is considered in isolation, without analyzing individual applications/services and without including the actual context and, moreover, it is limited to a few predefined user models, this does not enable power consumption optimization in the case of real-world mobile applications and services. Studies conducted by the authors mainly concern the prediction of energy efficiency of individual interface configurations; we think that subsequently, this knowledge is not adequately leveraged. In this respect our studies use power analysis, including the predicted power consumption, to optimize (using machine learning algorithms) the execution of individual services, which allows for the optimization of the entire mobile device. Conclusions from the analysis of the article described here were among our motivations to carry out our research and to develop a completely different, and innovative, approach to the topic of mobile device energy consumption optimization with the use of machine learning algorithms.

This research is a continuation of the results published in our papers [21, 22]. In [21], we present possibilities for using machine learning and the code offloading mechanism in the MCC concept, whereas in [22], we propose an innovative recommender system that allows the optimization of the selection of multimedia services (for converting photos and videos). The main goal in [22] was to optimize the service execution time. The system developed allows the selection of multimedia services offered by different providers locally on the mobile device or remotely in the cloud. In order to choose the service execution location, the concept of learning agents is used that utilizes various machine learning algorithms such as C4.5, Random Forest and Naïve Bayes. The tests conducted included the context associated with the type of network connection (LTE/HSPA/EDGE) and also various methods for converting photos and videos. The article focuses primarily on optimizing the service execution time; energy aspects have not been studied thoroughly. At the same time, the solution developed only covers a system for recommending choices with respect to a single type of service (multimedia conversion), which limits its versatility to some extent.

With respect to using the Mobile Cloud Computing concept, an important aspect is that the optimization of resources concerns not only the mobile device but the cloud as well [23] [24]. In [23], the authors demonstrate that it is possible to optimize resource usage in MCC by applying common patterns used in traditional cloud computing, whereas in [24], the authors propose optimization the power consumption in the data center based on neural network.

3 Supervised Learning Techniques in Context of Mobile Devices

Generally, supervised learning allows us to generate an approximation of the function $f : X \rightarrow C$, which assigns labels from set C to objects from set X . To generate knowledge, a supervised learning algorithm requires labeled examples that consist of pairs of f arguments and values. Let us assume that elements of X are described by a set of attributes $A = (a_1, a_2, \dots, a_n)$ where $a_i : X \rightarrow D_i$. Therefore $x^A = (a_1(x), a_2(x), \dots, a_n(x))$ is used

Table 1 Comparison of supervised learning algorithms applied in mobile devices

Learning algorithm	Learning speed	Hypothesis complexity	Readability to humans
NB	Fast	Low	Poor
C4.5	Medium	High	Good
Random forest	Slow	Very high	Poor

instead of x . If the size of C is small, like in this study, the learning is called classification, C is the set of classes, and h is called the classifier.

The supervised learning module obtains *Training data*, which is a set $\{(x^A, f(x))\}$, and generates the hypothesis h , which is stored in *Generated Knowledge*. The *Problem* is x^A , and the *Answer* is $h(x^A)$.

There are many supervised learning methods, which use various hypothesis representations and various methods of constructing hypotheses. Three of them are described below.

Naïve Bayes (NB) is a simple probabilistic classifier, which is a special case of a Bayesian network. Generally, a Bayesian network is a pair (G, P) where G is a structure graph and P is a set of local, conditional probability distributions between variables and their parents. In the case of NB, G is very simple: the class node c is the parent of every attribute node a_1, a_2, \dots, a_n . Learning is a process of calculating a priori probabilities $P(c)$ and conditional probabilities $P(a_i|c)$. The probability distribution of a class variable, for example $(a_1(x), a_2(x), \dots, a_n(x))$, is calculated using the following formula:

$$P(c(x) = c_j) = \frac{P(c_j) \prod_{i=1}^n P(a_i(x)|c_j)}{\prod_{i=1}^n P(a_i(x))} \quad (1)$$

C4.5 is a decision tree learning algorithm developed by Ross Quinlan [25]. The basic idea of learning is as follows: the tree is learned from examples recursively. If (almost) all examples in the training data belong to one class, the tree consisting of the leaf labeled by this class is returned. In the other case, the best attribute for the test in the root is chosen (using an entropy measure), training examples are divided according to the selected attribute values, and the procedure is called recursively (for every attribute test result with the rest of attributes and appropriate examples as parameters.)

The random forest algorithm builds an ensemble of decision trees. To reduce overfitting, trees are trained on subsets of the training set (selected at random) and using subsets of the attribute set (also selected at random) [26]. The decision is calculated by voting.

In the context of the application considered, these three algorithms may be compared as follows (see Table 1): NB is the fastest, but it is able to account for dependencies between attributes. The knowledge represented by probabilities is difficult to analyze. C4.5 learns more slowly, but the decision tree may represent any hypothesis. It is also readable for human experts. Random forest is the slowest algorithm but exhibits robust prediction results. Trees may represent complex hypotheses and ensembles make it possible to reduce overfitting. However, because of a large number of trees, the knowledge generated is also difficult to analyze.

In reinforcement learning, various techniques are used to prevent reaching a local optimum. The idea is to explore the solution space more thoroughly by choosing suboptimal actions from time to time (e.g. random or not performed in a given state yet). We decided

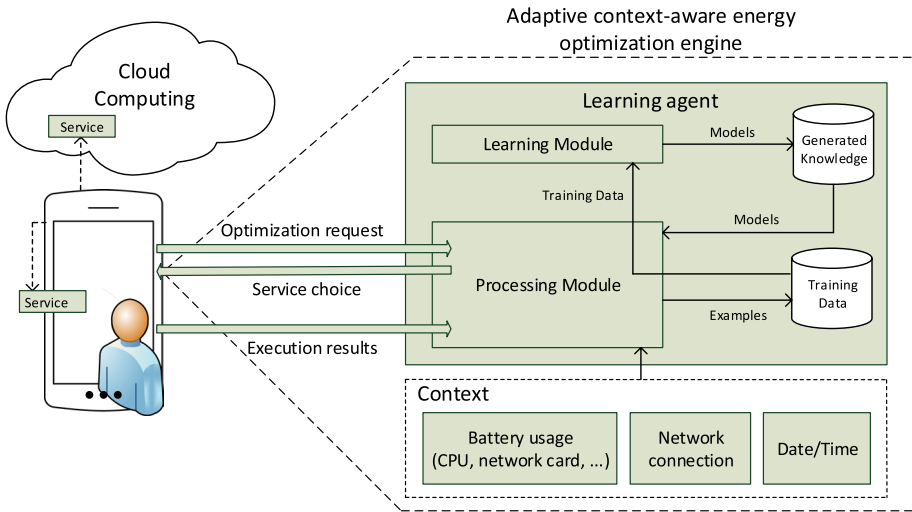


Fig. 1 Architecture of the adaptive service choice system

to use this technique for supervised learning too because experimental results show that exploration is useful.

One such method is ϵ -greedy where the agent selects the action that it believes has the best long-term effect with probability $1 - \epsilon$, and it chooses an action uniformly at random otherwise. $\epsilon \in (0, 1)$ is a tuning parameter. In our experiments it is constant, but in reinforcement learning it often decreases with time.

4 Adaptive Context-Aware Energy Optimization for Services

Let us define the learning agent (Fig. 1) as a tuple:

$$LA = (P, L, TD, K, T, C, R, A, D), \tag{2}$$

where P is the processing module, L is the learning module, TD is training data, K is knowledge learned, T is a set of computational tasks, C is a set of possible contexts (battery state, connection, date, etc.), R represents possible task execution results, A is a set of attributes, which are used to describe tasks, results and the context, and D is a set of decisions. The aim of the agent is to return decision $d \in D = \{d_1, d_2, \dots, d_{ns}\}$, which corresponds to engaging one of ns services.

Input data for processing module P is a pair $x = (t, c) \in T \times C$. This pair describes it with attributes from $O \subset A$, which yields $x^O = (o_1(x), o_2(x), \dots, o_n(x))$. Next, using the knowledge stored in K it selects $d \in D$, which has the minimum predicted cost. If K is empty, d is randomized.

The decision d is then applied and the task is run using the corresponding service (e.g. locally or in the mobile cloud). After the execution, the P module obtains execution results $r \in R$, which are described by $Res = \{r_1, r_2, \dots, r_m\} \subset A$ attributes (e.g. battery consumption $b(x, d)$, calculation time $ct(x, d)$).

The P module stores those results together with x^O and decision d in TD . Therefore the complete example stored in TD has the form

$$x^{AUD} = (o_1(x), o_2(x), \dots, o_n(x), r_1(x, d), r_2(x, d), \dots, r_m(x, d), d). \quad (3)$$

The knowledge used to predict R values is trained by the learning module L using supervised learning algorithms and TD , and stored in K . The form of the knowledge depends on the learning algorithm applied. For example, if C4.5 is used, the knowledge has the form of a decision tree with leaves representing values of predicted resource usage (time or battery) and other nodes represent tests on O attributes. In cases like this, when the learning algorithm performs classification instead of regression, the predicted attribute must be discretized. The number of bins has a high impact on accuracy. However, five bins turned out to be the optimal value in all our cases (see the tuning of parameters described in Sect. 5.3).

Using value predictions $r_i \in Res$, the processing module P rates its decisions $d \in D$ by calculating predicted expenses $e(x, d)$:

$$e(x, d) = \sum_{i=1}^m w_i * r_i(x, d), \quad (4)$$

where w_i are weights of the result r_i . By choosing the weights, one sets priorities for the criteria. As a result, the system is flexible and universal, because it may be adjusted to user requirements.

The P module selects the decision for which execution is predicted to be successful and the expense is predicted to be the lowest. To avoid a local optimum, the ϵ -greedy strategy is applied and a suboptimal decision is executed from time to time. The algorithm of the P module including the steps described above is presented in Algorithm 1.

```

1  Get task and the context  $x = (t, c) \in T \times C$ ;
   /* For example  $t = \text{OCR}$ ,  $c = \text{Wi-Fi}$  */
2  if  $K \neq \emptyset$  then
   /* Prediction models already learned */
3    foreach  $d \in D$  and  $r_i \in Res$  do
4      Predict results  $r_i(x, d)$  using models from  $K$ ;
      /* For all possible execution location predict time and energy
         consumption */
5    end
6    Select best decision:  $d = \arg \min_d \sum_{i=1}^m w_i * r_i(x, d)$  ;
      /*  $d$  represents the best execution location */
7     $rnd =$  random number between 0 and 1 ;
      /* Randomize to check if perform exploration */
8    if  $rnd < \epsilon$  then
9       $d =$  random decision ;
      /* Override  $d$  to perform exploration */
10   end
11 end
12 else
   /* There are no prediction models in the first round, therefore execution
      location is random */
13    $d =$  random decision ;
14 end
15 Execute task  $t$  using decision  $d$  ;
   /* For example execute OCR task in the cloud */
16 Observe results  $r \in R$  ;
   /* Measure time and energy consumption */
17 Store example
    $x^{A \cup D} = (o_1(x), o_2(x), \dots, o_n(x), r_1(x, d), r_2(x, d), r_m(x, d), d)$  in  $TD$ ;
   /* Task, context, time, energy consumption and execution location are stored in
      training data, from which prediction models will be learned */
18 if It is learning time then
19   Learn new  $K$  from  $TD$ :  $K = L(TD)$ ;
   /* Learn time and energy prediction models */
20 end
21 Goto 1;

```

Algorithm 1: Adaptation algorithm

5 Evaluation

In order to run experiments, software using the architecture developed was implemented, which makes it possible for a set of tests to be conducted and detailed results to be

Table 2 Types of tests for the Face Recognition service

Name	Resolution	Video frames	Duration	Comments
fr1	1280 × 720	29 fps	00:00:03	Single face
fr2	640 × 480	29 fps	00:00:31	Three different faces, rapid change in the face currently displayed
fr3	1280 × 720	30 fps	00:00:25	More than a dozen faces visible at the same time
fr4	406 × 720	30 fps	00:00:03	Single face
fr5	854 × 476	25 fps	00:00:13	Three different faces, rapid change in the face currently displayed

generated for a chosen configuration. The software developed can operate in two modes: for determining optimal classifier parameters (Hill Climbing) and for performing a set of tests. Both modes use the architecture developed, but in the first mode classifier parameters are modified on the go while the second mode uses predetermined values of classifier parameters.

All tests carried out consisted of multiple series and each series comprised a set of rounds. During each round, a set of tasks were performed (Face Recognition or OCR). After obtaining the result for the individual task, the time r_i and power consumption r_p were measured (i.e. costs of performing the task). After each round had been conducted, a classifier was built on the basis of the knowledge acquired from the previous and current rounds in the series in question. Those classifier parameters were established in the process of Hill Climbing. In the first round (reference round), the task execution location (cloud or mobile device) was selected randomly. The costs calculated on the basis of the results of this round were not used to build a classifier, but to calculate the penalty that was applied when the task ended in an error. Such a situation could occur e.g. when the network connection established to execute the cloud service was disrupted or terminated. Then, instead of the costs calculated, r_i values from the reference round corresponding to this task were used, multiplied by a constant factor equal to 1.5. After all rounds in the series had been conducted, all knowledge gathered was deleted. During the tests carried out, multiple series were conducted in order to obtain average values for individual rounds.

Two services were used during tests: Face Recognition and OCR. For the Face Recognition (FR) service, five types of tests were developed (shown in Table 2) for different input data (video stream parameters).

For the OCR service, five types of test were developed as well (shown in Table 3) for different input data (image parameters).

Three mobile devices were used during initial tests: the Lenovo Tab 2 A7-30D (1.3 GHz CPU, 1 GB RAM) tablet as well as the Samsung Galaxy Trend Plus (1.2 GHz CPU, 768 MB RAM) and HTC Desire 610 (1.2 GHz CPU, 1 GB RAM) mobile phones. All devices used the Android 4.4.2 operating system. The main experiment was conducted using only the Lenovo Tab 2 A7-30D device. In order to run remote tasks (Face Recognition and OCR) in the cloud, the AWS Lambda solution was used.

All experiments were executed using real-world Internet connections (Wi-Fi and HSDPA/HSUPA). Therefore we were not able to control connection quality and the results obtained suffer from relatively high variation. However, such conditions are similar to real-world applications in which connection quality may change.

The following attributes are used in experiments: O consists of nine attributes presented in Table 4. Eight of them describe a task t , one (*connectionType*) describes a context c .

Table 3 Types of tests for the OCR service

Name	Resolution	Format	Size of file	Comments
ocr1	640 × 480	PNG	22.8 KB	Printed black text, white background, large font
ocr2	720 × 255	JPEG	141 KB	Typescript, partly faded, newspaper article clipping
ocr3	3768 × 5256	JPEG	2.26 MB	Printed black text, white background, large font
ocr4	716 × 484	JPEG	73.3 KB	Black printed text with purple header, white background, low image quality
ocr5	306 × 306	JPEG	22.2 KB	Photo of a short text, gradient background (white and red), horizontal lines above and below the text

Table 4 Attributes describing tasks and the context

Name	Domain	Description
taskType	{fr, ocr}	Type of the task to execute
Size	\mathbb{N}	FR: number of frames, OCR: file size
Resolution	\mathbb{N}	Number of pixels (per frame for FR)
nrOfTemplates	\mathbb{N}	FR: number of faces to detect, OCR: N/A
handleRotations	Boolean	FR: value of HandleArbitraryRotations parameter, OCR: N/A
determineRotation	Boolean	FR: value of DetermineFaceRotationAngle parameter, OCR: N/A
resizeWidth	\mathbb{N}	FR: value of InternalResizeWidth parameter, OCR: N/A
Threshold	\mathbb{N}	FR: value of DetectionThreshold parameter, OCR: N/A
connectionType	{none, HSDPA/ HSUPA, Wi-Fi, unknown}	Internet connection type

Execution results R consist of two numeric attributes: *batteryUsage* and *timeUsage*. The set of decisions D has two values: *cloud* and *local*.

5.1 Power Measurement

The power consumption measurement (estimation) module was an important element used in power consumption tests in the process of learning and also an element allowing the classifier that predicted energy demand to learn. The choice of the method for measuring or estimating power was preceded by analysing and testing existing solutions. The tests were performed for Android 4.4.2, which was used by a large number of mobile devices at the time (in 2016). The method selected should work in real time (online), allow measurements for individual device components (CPU, wireless communication modules), and should have appropriate measurement resolution (less than 1% battery consumption). Owing to this resolution, it is possible to measure the power consumption of particular applications/services on mobile devices quite accurately.

The first study concerned power measurement methods on Android devices. The most commonly used measurement method is to use a public API (*BatteryManager*) that makes it possible to retrieve information about the current power status of the device. It uses a subscription mechanism, which prevents obtaining information about battery status regularly and continuous real-time measurements. At the same time, the maximum resolution of measurements is 1%, which may sometimes not be sufficient to measure the difference in power consumption between application/service launches in different contexts. It is possible to use an advanced API (via the *android.os.BatteryManager* class), which allows measurements with a resolution better than 1%, but this is feasible only for a few devices with the Summit SMB347 and MAX17050 battery charger integrated circuits, which are present in Nexus series devices (such as Google Nexus 6 and 9), so this is not a solution which could be widely used on a variety of Android devices. There is also the non-public *BatteryInfo* Android API, which makes it possible to obtain low-level information about power consumption. However, it requires *android.permission.BATTERY_STATS* permissions, which are reserved for applications built into the system and cannot be easily used by user applications. Another option for obtaining

data on power consumption is to use the *dumpsys batterystats* command and visualize the results in the Battery Stats and Battery Historian programs. This enables data to be obtained from system logs, including information on power consumption by the entire device as well as its individual components. The downside of this solution is that it does not work online and that it only shows the results on a PC. Another solution (Carat [27]) allows for the monitoring and analysis of data (on an external server) on power consumption from multiple mobile devices simultaneously and for detecting anomalies in the operation of individual applications. However, this solution does not allow local power measurement in real time (online). There are also closed applications available, such as Battery Doctor, Battery Saver 2017 and GSam Battery Monitor, which can be used for the monitoring and management of battery consumption on mobile devices. Still, due to the lack of their source codes or libraries, it was not possible to use them in the solution developed. An alternative to software solutions is the physical measurement of battery power consumption. This is the most accurate method, but it usually requires gaining access to mobile device internals and the use of additional measuring equipment. Moreover, such solutions only allow for the measurement of total power consumption of the device without obtaining any results for individual components such as CPU or wireless communication modules. Because of the nature of this measurement, this solution will never be widely used. Examples in this category are BattOr [28] (open-source), and the commercial Monsoon Mobile Device Power Monitor solution.

Our further research focused on the estimation of power consumption by mobile devices. Solutions using this method make it possible to obtain real-time results with a high measurement resolution. The most popular and widely used solution based on estimation is the PowerTutor program, which uses three basic energy characteristics of mobile device components. For a thorough analysis of the capabilities of this solution and its potential further use, we used the source code of this program to develop our own library for estimating power consumption. The solution developed allows for estimating power consumption of individual mobile device components (such as CPU and communication modules) with an error in the range of 1–5% [20], but it was designed for older devices and does not support new LTE wireless communication modules. Finally, the latest method for estimating mobile device power consumption uses the power profiles provided by device manufacturers. However, not all mobile devices have these profiles defined correctly which leads to problems with using this method. At the same time, no libraries have been developed that use this solution. In order to compare this method with PowerTutor, we conducted tests using the Lenovo Tab 2 A7-30D device. Preliminary results showed that for this device, there are no major differences between the two solutions when it comes to estimating the power consumption of the CPU and wireless communication modules.

The analysis of available solutions demonstrates that only estimation methods work on most devices and meet basic requirements, i.e. they work online, allow for measuring individual components of a mobile device and have the appropriate measurement resolution. While being aware of its limitations, we have decided to use our own library developed with the use of the PowerTutor source code. However, with the device used in tests, this solution allowed for a fairly accurate estimation of the power consumed by the CPU and wireless communication modules. In future research concerning new Android mobile devices, we are going to use a solution based on power profiles, which has been pre-tested by us. This will require developing a library for analysing the *power_profile.xml* system file and calculating the power consumption of individual components of a mobile device.




Fig. 2 Battery consumption while building a classifier

5.2 Power Consumption During Learning Process

In the first stage of proper research, the energy cost of building classifiers was measured (including the operation of the Weka library [29, 30]), which made it possible to assess whether it was not excessive in relation to the potential savings resulting from the use of the solution developed. Tests were performed for the three classifiers used (C4.5², Random Forest, Naïve Bayes) using the Lenovo Tab 2 A7-30D tablet. For conducting tests, artificial training data were generated, containing 1,000 random examples, which corresponds to executing 1,000 tasks using the system developed. Every classifier was tested with different percentages (25%, 50%, 75% and 100%)³ of the training data set used. Tests for specific percentage values were repeated 100 times and the final result was averaged. In Fig. 2, battery consumption (as a percentage) during the process of building individual classifiers is shown. The results demonstrate that power consumption for all classifiers is low (up to 0.4% of battery charge) and does not significantly affect the ability to carry out the tests of the services developed. In addition, the data set used for these tests was very large and in practice when it comes Face Recognition and OCR service tests, the amount of data that had to be processed by the classifier was much smaller (usually from 100 to 200 examples). The lowest power consumption is associated with the Naïve Bayes classifier and this is due to the fact that Naïve Bayes exhibits linear time complexity, which results in less stress on the device at the time of building the classifier (compared to the remaining classifiers used) and thus less energy expenditure.

² Weka's C4.5 implementation (J48) was used in the tests.

³ For example, a 25% test meant that to build the classifier, 250 examples were used from the input data.

Table 5 Optimization of parameters for the C4.5 classifier

Name	Initial value	Δ	Final value
Pruned	False	“Toggle”	False
ϵ	10.0	5.0	15.0
Bins (for class attributes)	5	2	5
minNumObj	3	1	3

Table 6 Optimization of parameters for the Naïve Bayes classifier

Name	Initial value	Δ	Final value
ϵ	10.0	5.0	20.0
Bins (for class attributes)	5	2	5

5.3 Tuning Classifier Parameters

In the next research stage, the Lenovo Tab 2 A7-30D mobile device was used, on which the Hill Climbing algorithm was run in order to determine the optimal parameters for individual classifiers. During the tests, the weights of adaptive algorithm were set as follows: $w_p = 50$ and $w_t = 50$, and for each classifier, 20 series of the algorithm were run, each series consisting of five rounds, and all series for a single set of parameters were repeated twice. Initial values of parameters were selected manually based on results of a number of optimization process executions during software development and testing. They appeared to be close to optimal values. However, one may start the tuning process from another starting point, taking into account that it may take longer. Generally, the Δ parameter should be as small as possible given the time complexity of the process. As concerns the number of bins, it equals two because we wanted odd values of this parameter to have a neutral value in the middle. This is not necessary, though. The Δ parameter for ϵ is set to five to limit the amount of computations. For each set of parameters, tasks were executed in various contexts—a single test round consisted of:

- Five tasks executed with the Wi-Fi connection available (9 Mb/s), including three Face Recognition tasks (fr3, fr4, fr5) and two OCR tasks (ocr4, ocr5);
- Five tasks with the HSDPA/HSUPA connection available, including three Face Recognition tasks (fr3, fr4, fr5) and two OCR tasks (ocr4, ocr5).

The result of this stage of research was the optimization of parameters for two classifiers: C4.5 (Table 5) where the ϵ parameter was changed (from 10 to 15) and Naïve Bayes (Table 6) where the same parameter changed from 10 to 20. Such a big change in the value of the Naïve Bayes classifier probably resulted from the fact that the decisions made by the classifier were sometimes suboptimal and a higher value of this factor contributed to better optimization.

For the Random Forest classifier (Table 7), the initial parameters proved to be optimal and did not require improvement.

Table 7 Optimization of parameters for the Random Forest classifier

Name	Initial value	Δ	Final value
ϵ	10.0	5.0	10.0
Bins (for class attributes)	5	2	5
numFeatures	3	1	3

5.4 Optimization

The final stage consisted of conducting the tests concerning the possibility of optimizing power consumption (and additionally execution time) using different classifiers. During the tests related to optimizing power consumption, the weights of the adaptive algorithm were set to $w_p = 90$ and $w_t = 10$. For the additional test related to optimizing execution time (Fig. 10), the weights were set to $w_p = 10$ and $w_t = 90$. The Lenovo Tab 2 A7-30D mobile device was used during test. Each test carried out consisted of 20 series, each series of nine rounds, and each round comprised:

- Seven tasks executed with the Wi-Fi connection available (9 Mb/s), including four tasks of the Face Recognition type (fr1, fr2, fr4, fr5) and three tasks of the OCR type (ocr1, ocr2, ocr4);
- Seven tasks executed with the HSDPA/HSUPA connection available, including four tasks of the Face Recognition type (fr1, fr2, fr4, fr5) and three tasks of the OCR type (ocr1, ocr2, ocr4).

Detailed test results for individual classifiers are presented in two graphs: for the optimization of power consumption and task execution location (mobile device or cloud computing). For the power consumption optimization graph, the result of a single round was the aggregate power consumption by all tasks executed in that round. The graph shows the average results of individual rounds conducted in all series; for each result, the standard deviation is marked. The t-Student test was also performed for each classifier and for averaged results from all series of the first and the last round. In the case of the graph showing the task execution location, the result of a single round is the number of tasks executed in a given location (locally on the mobile device or remotely using cloud computing). On the graph, each round is marked separately and contains the average number (from all series) of tasks executed in a particular location (locally/remotely).

Figure 3 shows the power consumption optimization graph for the C4.5 classifier. It can be seen that power consumption decreases in subsequent rounds until it begins to oscillate around a single value of 17,500 mJ. The result of the t-Student test for this classifier (the p-value) equals 0.000018, which means that average values for the first and last rounds are statistically significantly different.

Figure 4 shows the graph presenting the number of tasks executed locally on the mobile device and remotely in the cloud for the C4.5 classifier. It can be noticed that the algorithm using this classifier sends more and more tasks to the cloud over time (in successive rounds), reducing power consumption on the mobile device.

Figure 5 shows the power consumption optimization graph for the Random Forest classifier. It can be seen that power consumption decreases in subsequent rounds until it reaches a value of about 18,000 mJ. The result of the t-Student test for this classifier (the




Fig. 3 Energy optimization for the C4.5 classifier




Fig. 4 Number of tasks performed in the cloud and on the mobile device for the C4.5 classifier

p -value) equals 0.00000017, which means that average values for the first and last rounds are statistically significantly different.

Figure 6 shows the graph presenting the number of tasks executed locally on the mobile device and remotely in the cloud for the Random Forest classifier. It can be noted that the algorithm that uses this classifier, similarly as in the case of C4.5, sends more tasks to the cloud, reducing power consumption on the mobile device.




Fig. 5 Energy optimization for the Random Forest classifier




Fig. 6 Number of tasks performed in the cloud and on the mobile device for the Random Forest classifier

Figure 7 shows the power consumption optimization graph for the Naïve Bayes classifier. It can be seen that power consumption between the first and last rounds does decrease, but it is not a steady or large reduction. The result of the t-Student test for this classifier (the p-value) equals 0.1, which means that average values for the first and last rounds are not statistically significantly different.




Fig. 7 Energy optimization for the Naïve Bayes classifier




Fig. 8 Number of tasks performed in the cloud and on the mobile device for the Naïve Bayes classifier

Figure 8 shows a graph presenting the number of tasks executed locally on the mobile device and remotely in the cloud for the Naïve Bayes classifier. It can be noticed that the algorithm using this classifier, similarly to the previous tests, sends more tasks to the cloud; however, it does not allow a significant power consumption optimization to be achieved. This might be related to the high value of the random factor, which was the result of running the Hill Climbing algorithm.




Fig. 9 Comparison of power consumption optimization for all three classifiers (C4.5, Random Forest and Naïve Bayes) and for services performed without using machine learning methods (exclusively locally on the mobile device and exclusively in the cloud)

Figure 9 shows a comparison of power consumption optimization results for all three classifiers and for services performed without using machine learning methods (exclusively locally on the mobile device and exclusively in the cloud). It can be noticed that classifiers based on decision trees (C4.5 and Random Forest) perform much better than the Naïve Bayes classifier. They achieve almost the same optimization levels and the result of the t-Student test for both classifiers (the p-value) equals 0.3741, which means that average values for the last round for the C4.5 and Random Forest classifiers are not statistically significantly different. However, the Random Forest classifier achieves the final power consumption stage faster. The worst result was achieved by the Naïve Bayes classifier. The result of the t-Student test for the Naïve Bayes and C4.5 classifiers (the p-value) equals 0.0021, which means that average values for the last round for those classifiers are statistically significantly different. In cases where the service was executed in a single location (in the cloud or locally), the results were worse than in cases where classifiers and machine learning were used. However, the result for running the service exclusively in the cloud was only slightly worse than for the Naïve Bayes classifier.

In order to check whether the algorithm developed allows for the optimization of other parameters, task execution time optimization tests were carried out (Fig. 10) for the C4.5, Random Forest and Naïve Bayes classifiers and for services performed without using machine learning methods (exclusively locally on the mobile device and exclusively in the cloud). For all the classifiers tested, task execution time between the first and last rounds decreased. Similar to the power consumption optimization tests, classifiers using decision trees performed much better than the Naïve Bayes classifier when it came to optimizing task execution time. The result of the t-Student test for average values of the first and last rounds of Naïve Bayes classifier tests amounted to 0.3, which means that there is no statistically significant improvement in task execution time results for that classifier. For the C4.5 (the p-value in the t-Student test equals 0.00019) and Random Forest (the p-value in

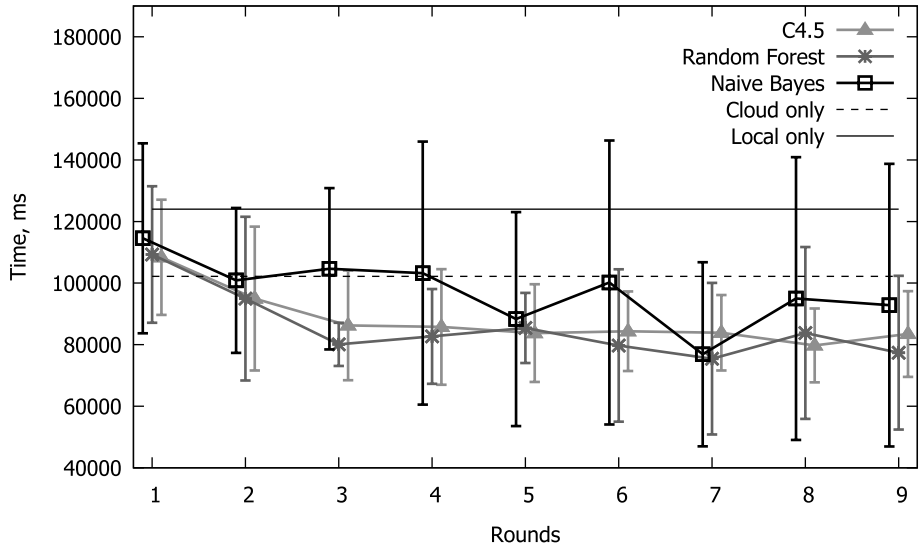


Fig. 10 Comparison of execution time optimization for all three classifiers (C4.5, Random Forest and Naïve Bayes) and for services performed without using machine learning methods (exclusively locally on the mobile device and exclusively in the cloud)

the t-Student test equals 0.0026) classifiers, the decrease in task execution time between the first and last rounds was statistically significant. All classifiers achieved almost the same level of optimization. Results of t-Student tests for these classifiers were as follows: 0.3471 (C4.5/Random Forest), 0.3880 (C4.5/Naïve Bayes) and 0.1939 (Random Forest/Naïve Bayes) which means that average values for the last round for all classifiers are not statistically significantly different. In cases where the service was executed in a single location (in the cloud or locally), the results were significantly worse than in cases where classifiers and machine learning were used.

Numbers of tasks executed in both locations for various contexts (HSDPA/HSUPA and Wi-Fi connections) are presented in Figs. 11, 12, 13, 14, 15 and 16. For all learning algorithms and contexts, the number of local executions exhibits a downward trend, while the number of cloud executions shows an upward one. The difference between local and cloud executions is larger for the Wi-Fi connection than for HSDPA/HSUPA because transfer speed is higher, the connection is more stable, and data transfer becomes cost-effective for a larger number of tasks. This is particularly noticeable for the C4.5 and Random Forest algorithms, which are more accurate than Naïve Bayes.

In order to compare our software with already existing solutions, we analyzed various Mobile Cloud Computing solutions. Many of those (such as MALMOS [31], COMET [32] and COSMOS [33]) do not take into account energy aspects at all in their operation. Only a few (such as AIOLOS [34], CACTSE [35], Cuckoo [36], EMCO [37], IC-Cloud [38], MAUI [39] and ThinkAir [40]) account for energy aspects and only the IC-Cloud solution uses machine learning algorithms to optimize the operation of applications/services. However, almost all of the solutions analyzed (including IC-Cloud) are not being developed any further or there is no access to their source codes. It was only




Fig. 11 Number of tasks performed in the cloud and on the mobile device for the Naïve Bayes classifier and HSDPA/HSUPA connection




Fig. 12 Number of tasks performed in the cloud and on the mobile device for the Naïve Bayes classifier and Wi-Fi connection

possible to find source codes for two solutions: AIOLOS⁴ and Cuckoo⁵. Unfortunately, both of these solutions use old software development kit versions (such as Eclipse and

⁴ Code repository—<https://github.com/ibcn-cloudlet/aiolos>.

⁵ Code repository—<https://github.com/interdroid/cuckoo-library>.




Fig. 13 Number of tasks performed in the cloud and on the mobile device for the C4.5 classifier and HSDPA/HSUPA connection




Fig. 14 Number of tasks performed in the cloud and on the mobile device for the C4.5 classifier and Wi-Fi connection

the ADT plugin instead of Android Studio). In the case of AIOLOS, we were able to configure and build a sample project, but when running the sample (using the Androsgui plugin), the application closed and reported an error. The code was analyzed, but it was not possible to determine what caused the error. For the second solution—Cuckoo, it was possible to run the sample application. However, comparing this solution with our




Fig. 15 Number of tasks performed in the cloud and on the mobile device for the Random Forest classifier and HSDPA/HSUPA connection




Fig. 16 Number of tasks performed in the cloud and on the mobile device for the Random Forest classifier and Wi-Fi connection

system proved difficult due to the fact that Cuckoo lacked machine learning mechanisms and used completely different solutions in terms of the computing cloud—the server ran on an EC2 instance and it was not possible to use the AWS Lambda service.

6 Conclusions

Multiple studies have been carried out recently concerning the possibilities of reducing the power consumption of mobile devices, primarily in the area of optimizing the functioning of their screens, processors or wireless communication modules. However, our analysis concerning existing solutions revealed that there are no systems that use machine learning for optimizing power consumption on mobile devices using MCC. In this article, we present an original concept for an adaptive system enabling the optimization of mobile device power consumption and at the same time taking into account the context of the device's operation. The use of machine learning algorithms in our solution allowed for the optimization of power consumption and reducing the service/application execution time on the mobile device.

Two services were used in the tests: Face Recognition and OCR. They were implemented on a mobile device with the Android operating system, with the ability to run them in the AWS Lambda compute cloud. The experiments carried out demonstrated that the adaptation algorithm developed made it possible to reduce power consumption for the C4.5 classifier by 41%, for the Random Forest classifier by 48% and for the Naïve Bayes classifier by 17%. Despite the relatively large standard deviation for the Naïve Bayes classifier, tests showed that the solutions developed can significantly reduce the power consumption of mobile devices. At the same time, we examined the possibility of optimizing the execution time of services/applications, and the results obtained—execution time being reduced by 23% (C4.5), 29% (Random Forest) and 19% (Naïve Bayes)—also demonstrated the effectiveness of the solution developed. During the tests, we also found out that the power consumption of the system developed during the learning process itself is very low and does not significantly affect the operation of the mobile device.

Our research shows that it is possible to create a system using adaptive algorithms based on machine learning that enables efficient learning and making increasingly better decisions. This allows for power consumption optimization as well as for reducing the service execution time on mobile devices. In further tests, we would like to address, among other things, the possibility of using distributed learning. Local models would be learned on mobile devices like it is the case now; however, part of the data collected would also be sent to the cloud where global models could be learned from aggregated data provided by many users. Such global models could subsequently be sent to mobile devices and allow them to use global knowledge locally.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Nawrocki, P., & Sliwa, A. (2016). Quality of experience in the context of mobile applications. *Computer Science*, 17(3), 371.
2. Noor, T. H., Zeadally, S., Alfazi, A., & Sheng, Q. Z. (2018). Mobile cloud computing: Challenges and future research directions. *Journal of Network and Computer Applications*, 115, 70–85.
3. Malhotra, A., Dhurandher, S. K., Gupta, M., & Kumar, B. (2018). Emcloud: A hierarchical volunteer cloud with explicit mobile devices. *International Journal of Communication Systems*, 31(17), e3812.
4. Sutton, R., & Barto, A. (1998). *Reinforcement Learning: An Introduction (Adaptive Computation and Machine Learning)*. Cambridge: The MIT Press.
5. Berral, J. L., Goiri, Í., Nou, R., Julià, F., Guitart, J., Gavaldà, R., & Torres, J. (2010). Towards energy-aware scheduling in data centers using machine learning. In *Proceedings of the 1st international conference on energy-efficient computing and networking. e-Energy '10* (pp. 215–224). ACM: New York, NY
6. Sniezynski, B. (2015). A strategy learning model for autonomous agents based on classification. *International Journal of Applied Mathematics and Computer Science*, 25(3), 471–482.
7. Berral, J.L., Goiri, Í., Nou, R., Julià, F., Fitó, J.O., Guitart, J., Gavaldà, R., & Torres, J. (2012). Toward energy-aware scheduling using machine learning. In *Energy-efficient distributed computing systems* (pp. 215–244). Wiley
8. Flinn, J., & Satyanarayanan, M. (1999). Energy-aware adaptation for mobile applications. *SIGOPS Operating System on Review*, 33(5), 48–63.
9. Vallina-Rodriguez, N., & Crowcroft, J. (2013). Energy management techniques in modern mobile handsets. *IEEE Communications Surveys and Tutorials*, 15(1), 179–198.
10. Bhatia, G., Mahajan, R., & Khatri, S. K. (2017). A study for improving energy efficiency in mobile devices. In *International Conference on Infocom Technologies and Unmanned Systems (Trends and Future Directions) (ICTUS)* (pp. 588–592).
11. Meneses-Viveros, A., Hernández-Rubio, E., Mendoza, S., Rodríguez, J., & Quintos, A. B. M. (2018). Energy saving strategies in the design of mobile device applications. *Sustainable Computing: Informatics and Systems*, 19, 86–95.
12. Kumar, K., & Lu, Y. H. (2010). Cloud computing for mobile users: Can offloading computation save energy? *Computer*, 43, 51–56.
13. Zhang, Y., Niyato, D., & Wang, P. (2015). Offloading in mobile cloudlet systems with intermittent connectivity. *IEEE Transactions on Mobile Computing*, 14(12), 2516–2529.
14. De, D., Mukherjee, A., & Roy, D. G. (2020). Power and delay efficient multilevel offloading strategies for mobile cloud computing. *Wireless Personal Communications*, 112, 2159–2186.
15. Moldovan, A. N., Weibelzahl, S., & Hava Muntean, C. (2014). Energy-aware mobile learning: Opportunities and challenges. *IEEE Communications Surveys and Tutorials*, 16(1), 234–265.
16. Yang, S., Kwon, D., Yi, H., Cho, Y., Kwon, Y., & Paek, Y. (2014). Techniques to minimize state transfer costs for dynamic execution offloading in mobile cloud computing. *IEEE Transactions on Mobile Computing*, 13(11), 2648–2660.
17. Chun, B.G., Ihm, S., Maniatis, P., Naik, M., & Patti, A. (2011). Clonecloud: Elastic execution between mobile device and cloud. In *Proceedings of the 6th Conference on Computer Systems. EuroSys '11* (pp. 301–304). New York: ACM
18. Hu, T., & Fei, Y. (2010). Qelar: A machine-learning-based adaptive routing protocol for energy-efficient and lifetime-extended underwater sensor networks. *IEEE Transactions on Mobile Computing*, 9(6), 796–809.
19. Donohoo, B. K., Ohlsen, C., Pasricha, S., Xiang, Y., & Anderson, C. (2014). Context-aware energy enhancements for smart mobile devices. *IEEE Transactions on Mobile Computing*, 13(8), 1720–1732.
20. Zhang, L., Tiwana, B., Qian, Z., Wang, Z., Dick, R. P., Mao, Z. M., & Yang, L. (2010). Accurate online power estimation and automatic battery behavior based power model generation for smartphones. In *Proceedings of the 8th IEEE/ACM/IFIP International Conference on Hardware/Software Codesign and System Synthesis. CODES/ISSS '10* (pp. 105–114). New York, NY: ACM.
21. Nawrocki, P., Sniezynski, B., & Slojewski, H. (2019). Adaptable mobile cloud computing environment with code transfer based on machine learning. *Pervasive and Mobile Computing*, 57, 49–63.
22. Nawrocki, P., Sniezynski, B., & Czyzewski, J. (2016). Learning agent for a service-oriented context-aware recommender system in a heterogeneous environment. *Computing and Informatics*, 35(5), 1.
23. Nawrocki, P., & Reszelewski, W. (2017). Resource usage optimization in mobile cloud computing. *Computer Communications*, 99, 1–12.
24. Akki, P., & Vijayarajan, V. (2020). Energy efficient resource scheduling using optimization based neural network in mobile cloud computing. *Wireless Personal Communications*.
25. Quinlan, J. (1993). *C4.5: Programs for Machine Learning*. Burlington: Morgan Kaufmann.

26. Breiman, L. (2001). Random forests. *Machine Learning*, 45(1), 5–32.
27. Oliner, A. J., Iyer, A. P., Stoica, I., Lagerspetz, E., & Tarkoma, S. (2013). Carat: Collaborative energy diagnosis for mobile devices. In *Proceedings of the 11th ACM Conference on Embedded Networked Sensor Systems. SenSys '13* (10:1–10:14) New York, NY: ACM.
28. Schulman, A., Schmid, T., Dutta, P., & Spring, N. (2011). Demo: Phone power monitoring with batter. In *ACM Mobicom*.
29. Hall, M., Frank, E., Holmes, G., Pfahringer, B., Reutemann, P., & Witten, I. H. (2009). The weka data mining software: An update. *SIGKDD Exploration Newsletter*, 11(1), 10–18.
30. Azeez, N. A., Asuzu, O. J., Misra, S., Adewumi, A., Ahuja, R., & Maskeliunas, R. (2018). *Comparative Evaluation of Machine Learning Algorithms for Network Intrusion Detection Using Weka* (pp. 195–208). Singapore: Springer.
31. Eom, H., Figueiredo, R., Cai, H., Zhang, Y., & Huang, G. (2015). Malmos: Machine learning-based mobile offloading scheduler with online training. In *Proceedings of the 2015 3rd IEEE International Conference on Mobile Cloud Computing, Services, and Engineering. MOBILECLOUD '15* (pp. 51–60). Washington, DC: IEEE Computer Society.
32. Gordon, M. S., Jamshidi, D. A., Mahlke, S., Mao, Z. M., & Chen, X. (2012). Comet: Code offload by migrating execution transparently. In *Proceedings of the 10th USENIX Conference on Operating Systems Design and Implementation. OSDI '12* (93–106). Berkeley, CA: USENIX Association.
33. Shi, C., Habak, K., Pandurangan, P., Ammar, M., Naik, M., & Zegura, E. (2014). Cosmos: Computation offloading as a service for mobile devices. In *Proceedings of the 15th ACM International Symposium on Mobile Ad Hoc Networking and Computing. MobiHoc '14* (pp. 287–296). New York, NY: ACM.
34. Verbelen, T., Simoens, P., De Turck, F., & Dhoedt, B. (2012). Aiolos: Middleware for improving mobile application performance through cyber foraging. *Journal of System Software*, 85(11), 2629–2639.
35. Qing, W., Zheng, H., Ming, W., & Haifeng, L. (2013). Cactse: Cloudlet aided cooperative terminals service environment for mobile proximity content delivery. *China Communications*, 10(6), 47–59.
36. Kemp, R., Palmer, N., Kielmann, T., & Bal, H. (2012). *Cuckoo: A computation offloading framework for smartphones* (pp. 59–79). Berlin, Heidelberg: Springer.
37. Flores, H., & Srirama, S. (2013). Adaptive code offloading for mobile cloud applications: Exploiting fuzzy sets and evidence-based learning. In *Proceeding of the Fourth ACM Workshop on Mobile Cloud Computing and Services. MCS '13* (pp. 9–16). New York, NY: ACM.
38. Shi, C., Pandurangan, P., Ni, K., Yang, J., Ammar, M., Naik, M., & Zegura, E. (2013). *ic-cloud: Computation offloading to an intermittently-connected cloud*. Technical report
39. Cuervo, E., Balasubramanian, A., Cho, D. K., Wolman, A., Saroiu, S., Chandra, R., & Bahl, P. (2010). Maui: Making smartphones last longer with code offload. In *Proceedings of the 8th International Conference on Mobile Systems, Applications, and Services. MobiSys '10* (pp. 49–62). New York, NY: ACM.
40. Kosta, S., Aucinas, A., Hui, P., Mortier, R., & Zhang, X. (March 2012). Thinkair: Dynamic resource allocation and parallel execution in the cloud for mobile code offloading. In *Proceedings IEEE INFOCOM* (pp. 945–953).

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Piotr Nawrocki Ph.D., is an Associate Professor in the Department of Computer Science at the AGH University of Science and Technology, Krakow, Poland. His research interests include distributed systems, computer networks, mobile systems, mobile cloud computing, Internet of Things and service-oriented architectures. He has participated in several EU research projects including MECCANO, 6WINIT and UniversAAL. He is a member of the Polish Information Processing Society (PTI).



Bartłomiej Sniezynski received his Ph.D. degree in Computer Science in 2004 from the AGH University of Science and Technology, Krakow, Poland. In 2004, he worked as a Postdoctoral Fellow under the supervision of Professor R. S. Michalski at the Machine Learning and Inference Laboratory, George Mason University, Fairfax, VA, USA. Currently, he is an Associate Professor in the Department of Computer Science at AGH. His research interests include machine learning, multi-agent systems and knowledge engineering. He is a member of the Polish Information Processing Society (PTI) and the Polish Artificial Intelligence Society (PSSI).