6-2014

The Case for Human-Centric Personal Analytics

Youngki LEE
Singapore Management University, YOUNGKILEE@smu.edu.sg

Rajesh Krishna BALAN
Singapore Management University, rajesh@smu.edu.sg

Citation

LEE, Youngki and BALAN, Rajesh Krishna. The Case for Human-Centric Personal Analytics. (2014). *WPA'14: Proceedings of the 2014 ACM Workshop on Physical Analytics: June 16, 2014, Bretton Woods, NH*. 25-29.
Available at: https://ink.library.smu.edu.sg/sis_research/2657

Published in WPA'14: Proceedings of the 2014 ACM Workshop on Physical Analytics: June 16, 2014, Bretton Woods, NH. 25-29. New York: ACM.
http://dx.doi.org/10.1145/2611264.2611267

The Case for Human-Centric Personal Analytics

Youngki Lee and Rajesh Krishna Balan
School of Information Systems, Singapore Management University
{youngkilee,rajesh}@smu.edu.sg

ABSTRACT

The rich context provided by smartphones has enabled many new context-aware applications. However, these applications still need to provide their own mechanisms to interpret low-level sensing data and generate high-level user states. In this paper, we propose the idea of building a personal analytics (PA) layer that will use inputs from multiple lower layer sources, such as sensor data (accelerometers, gyroscopes, etc.), phone data (call logs, application activity, etc.), and online sources (Twitter, Facebook posts, etc.) to generate high-level user contextual states (such as emotions, preferences, and engagements). Developers can then use the PA layer to easily build a new set of interesting and compelling applications. We describe several scenarios enabled by this new layer and present a proposed software architecture. We end with a description of some of the key research challenges that need to be solved to achieve this goal.

1. INTRODUCTION

Mobile computing has allowed unprecedented access to deep and enriched human contextual information, for example, mobility, activity, and interactions. This access is possible through smartphones that generate rich data sources, including application usage (e.g., social network postings, search history, and call records) and physical sensing data (e.g., location, activity). Recent advances in wearable devices are increasing this capability by providing health measurements as well as higher fidelity sensor data. Moreover, timely and real-time intervention based on user contexts and history is also now possible, as people usually carry their smartphones all the time.

This powerful capability has created new types of mobile applications. For example, persuasive well-being applications, that gently persuade users to change their undesirable life style habits based on *quantified self*, i.e., objective figures representing user behavioural patterns, are now fairly common. Overall, we are closer than ever to achieving the holy grail of context-aware mobile applications — where the applications automatically adapt based on what the user is doing, what they feel like doing, and what they should be doing. However, even with these new advances and decades of research efforts, it is still extremely challenging to achieve a usable level of context-awareness across a heterogeneous set of users and devices.

In this paper, we aim to get closer to this goal by proposing a new research direction known as *human-centric personal analytics*. We will first describe what this is and then describe the challenges to move the current state-of-the-art capabilities from *low-level physical sensing* to *high-level personal analytics*.

Human-centric personal analytics has two unique characteristics from current methodologies; First, it focuses on capturing high-level *human-centric contexts* (*intention, engagement, emotion, attention, fatigue, anxiety, depression, distractibility, mindfulness, etc.*) and not just low-level physical contexts such as location and activity. This was based on our observation and experience that context-aware applications are forced to infer user states from low-level inputs (such as location and activity) and usually do a bad job of this. Instead, it would be better if these low level inputs were already combined (with online profile data and other data sources), in smart ways, to provide applications with user states (such as intention or emotion) that are more useful for intervention purposes. For example, a future context-aware advertisement will not only target users based on their location and activity, but also more precisely target people who has buying intention to purchase certain products. Moreover, discount coupons can be sent out when the user is bored or has free attentive capacity.

Second, a flexible combination of real-time and historical data should be available for querying as this allows the discovery of deeper insights. For example to send out effective coffee shop coupons, the application will need a historic understanding of the user's coffee drinking behaviour. For example, to figure out if she even likes coffee and if she does, what time(s) she usually drinks coffee, what brands she likes, etc. The real-time capability will then be used to adjust the historical preferences to match the current situation. For example, if she has just passed by three coffee shops that she normally stops in, this might be a bad time to suggest having a coffee.

To enable human-centric personal analytics, we propose building a new mobile application service called the *Personal Analytics* (PA) Engine, which can be used by applications to *accurately*, *efficiently*, and *effectively* discover the human-centric contexts (HCC) of smartphone users.

HCCs are generated from inputs from multiple sources such as (1) multiple lower level sensing interfaces (such as accelerometer, gyroscope, and other phone sensors input) as well as (2) inputs from user context interfaces from both smartphones (such as call logs, browser history) and online sources (such as Facebook and Twitter). The PA engine will need to combine these raw data streams into a coherent consistent form that allows long-term and real-time monitoring of a variety of HCCs. Also, it must expose software APIs and runtime library that will allow external application developers, companies, and researchers to build their own novel and innovative user-centric applications that use HCCs to provide enhanced user experiences.

Recently, there have been several early efforts to detect several useful HCCs. For example, EmotionSense [12] and MoodScope [9] attempted to detect emotion of mobile users using diverse phone usage and sensor data. Also, StressSense [10] detects people's stress under interview situations using speech analysis based on recorded sound data. These early results showed the feasibility of smartphone-based human-centric personal analytics. However, these studies only targeted a small set of HCCs (mood and stress) in isolation and omitted many other attributes such as intention and mindfulness. In this paper, we are proposing a larger framework that will allow applications to access a broader range of interdependent HCCs with minimal programming effort. This is part of a larger effort, as part of the LiveLabs [2] testbed, to understand complex mobile user behaviours and actions.

In the rest of this paper, we will first discuss useful application scenarios that motivate the design of our PA engine. We then discuss our current thoughts on the system design and research issues involved in building the engine. We end with a brief description of our current status and future plans.

2. MOTIVATING SCENARIOS

Scenario 1: *CanICall*. John and Michelle are married to each other and work gruelling 8 a.m. to 8 p.m. jobs – John in the service sector and Michelle in the financial sector. They both were recently promoted which further reduced the time they spend with each other. They would like to talk more to each other during the day, even just shortly, but find it hard as they can't determine when the other person can be interrupted. To overcome this, they decide to install the *CanICall* application which uses the PA engine to determine if the user can be interrupted (by checking both contextual (calendar etc.) and sensing (ambient sound, light levels, etc.) inputs). If both parties can be interrupted and have not interacted with each other in some time, *CanICall* will prompt them to interact with each other. *CanICall* also suggests the most appropriate communication mode (call, SMS, gTalk, WhatsApp, Facebook poke, etc.) that is easiest to send and receive, at that point in time, by both parties. John and Michelle use *CanICall* for a week and realise that they are much happier as they communicate more often with no impact on work performance.

Scenario 2: *BoreMeter*. Ryan is attending his undergraduate mathematics class at the university. During the math lecture, Ryan gets bored and picks up his phone to talk to friends using WhatsApp. In the background, the *BoreMeter* application (that Ryan installed to help focus on his classes) is monitoring Ryans behaviour, via the PA engine, and discovers, through a combination of calendar, social media, and sensing inputs, that Ryan is in Math class and not engaged in the lecture. Immediately *BoreMeter* decides to stimulate Ryans interest in the lecture by launching a trivia game about the topic of the lecture. In addition, this incident is anonymously reported to the lecturer′s version of *BoreMeter* on his phone so that he can take corrective action if too many students start getting bored.

Scenario 3: *HappyShopper*. Vignesh enters a mall and starts looking for a gift for his wife. The *HappyShopper* application on his phone uses the PA engine to monitor Vigneshs reaction after he enters and leaves various types of shops. It notices that he appears noticeably happier (based on his facial expressions, gait, and comments sent via SMS to his wife) when entering and leaving dress shops as compared to shops selling other types of clothing (such as swimwear) and items (such as bags and shoes). *HappyShopper* thus decides to suggest to Vignesh various dress shops both in that mall and nearby that it believes will satisfy him the most (and also achieve the goal of buying his wife a nice present). Vignesh follows these directions and spends a very pleasant afternoon gift shopping.

Scenario 4: *WhoToHelp*. Mark is a sales manager at a large airport duty free shop. He has a limited number of sales staff to help customers. He would prefer using his staff to help serious buyers and not the majority of shopper who are just browsing to kill time between flights. To achieve this, he uses the *WhoToHelp* application, which discreetly tells his staff which customers are most likely to make purchases. The application uses the PA engine to determine a customer's buying intention (based on their walking patters, stay time inside the shop, and trajectory in the airport) and what products are most appealing to them (from available social network information and pause durations in front of various products). Mark uses *WhoToHelp* for a month and realises that the store's revenue and satisfaction ratings went up as his sales staff were able to increase their customer engagement capabilities even with limited human resource.

In all these scenarios, understanding the users′ HCCs (e.g., attention, boredom, happiness, engagement) are essential for building the various applications that can provide the proactive and useful interventions needed for each common day-to-day scenario.

3. PA ENGINE OVERVIEW

The key goals of the PA engine is to support the scenarios stated earlier by 1) providing accurate, low latency, en-

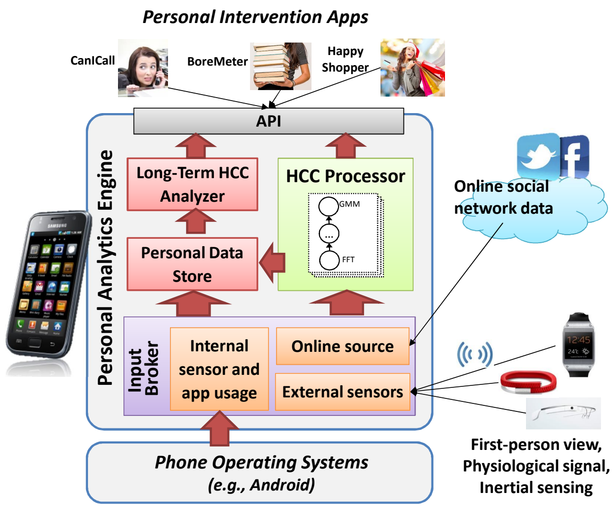

Figure 1: Architecture Overview of PA Engine

ergy efficient HCCs from sensors available in commodity smartphones and upcoming soon-to-be available wearable devices, and 2) provide a set of new APIs to allow real-time and historical HCC queries that can be readily used by application developers. In the rest of this section, we briefly introduce the four key technical components of our proposed system (shown in Figure 1).

Concurrent Sensor Optimiser: To provide *scalable*, *energy-efficient*, and *accurate* collection of lower level contextual data. To enable HCC-based inferencing, the system must collect data from a number of lower level sources, such as phone sensors, phone usage data, wearable devices, and online social network sources. We describe some of the challenges in building this component in Section 4.1.

HCC Processor: To provide *accurate* and *energy-efficient* inference of HCCs. This component must fuse the various lower level contextual data collected by the sensor optimiser to extract the user-specific HCCs. There have been recent efforts to detect several types of HCCs such as mood [9], stress [10], and emotion [12]. This component will extend those efforts to provide more accurate and expansive inferencing capabilities. We describe some of the challenges in building this component in Section 4.2.

Personal Data Store & Long-term HCC Analyser: To provide *privacy-aware* and *scalable* data storage coupled with long-term personal analytics capabilities. These two components work together to allow applications to perform meaningful long-term HCC analysis so as to detect longer-term patterns that can be used as triggers for real-time interventions. For example, the “BoreMeter” application could detect that afternoon classes tend to induce more boredom and prepare accordingly. We describe some of the challenges in building this component in Section 4.3.

API Abstraction and Runtime Engine. This final component will focus on integrating the above-mentioned basic technology blocks and providing a unified API that allows the development and deployment of future context-aware applications. The key challenge will be to provide intuitive, easy-to-use, yet expressive APIs that third-party application developers can readily use to create a wide variety of personal applications (such as those described in the scenarios) without requiring the developers to worry about collecting, storing, and processing lower level sensing information. Developers will also not need to worry about the accuracy, energy efficiency, or scalability of the various HCCs needed by their applications. Abstracting a clean set of usable and expressive API is not only a challenging design problem, but is also a basic necessity for the PA layer to have a larger impact beyond just in-lab developed research applications.

4. TECHNICAL CHALLENGES

In this section, we will present some of the key technical challenges in building the PA layer.

4.1 Collection of Primitive Data

Collecting a variety of primitive inputs (e.g., sensor data, phone data, and online data) is required to infer HCCs and perform long-term personal analytics. However, this collection raises various technical challenges:

Energy-efficient collection of sensor data. An obvious but important challenge is to collect this data in an energy efficient way. This is frequently the top priority as many mobile users are highly sensitive to changes in their phone's battery drain (even more so than privacy concerns from our experience!). As such, we need to achieve our data collection with just small battery drain (at most 5-10% only).

On the other hand, to enable diverse HCC-enabled applications and deep personal analytics, we need to collect various and fine-grained sensor data (e.g., regular WiFi scans with high frequency inertial sensing data to infer micro movements inside a store). For some applications, we might also require data from external wearable devices (e.g., heart rate from smart bands to infer acute stress, and video data from smart glasses to identify level of attention) which increases the overall power consumption.

There has been significant body of prior work to optimise continuous sensing applications, but they mostly focus on only a single application (with just a few sensing modalities) at one time. Some prior work such as ACE [11], KOBE [4], and MobiCon [8] provide generic frameworks to optimise the power consumption of multiple applications. However, it is still not a solved problem and the reality is much more complicated when we factor in device and application heterogeneity (e.g., where different phones have very different power consumptions for the same sensors).

Managing multiple applications also raises a unique challenges of managing the *shared power usage* across multiple sensing and processing tasks. The key idea is that new sensing tasks can benefit from already running tasks as many required components (e.g., sensors, CPU, memory, and flash disk) are already in power-on states — if the components are used smartly. Our prior measurement study [1] showed that this type of power sharing is possible as the power consumption does not scale linearly with the number of sensors used. However, a key research challenge is identifying the correct methods and ways to allow applications to share resources while still achieving all required accuracy and latency requirements.

Empowering new sensing devices. For accurate inference of diverse HCCs, it might not be sufficient to leverage data collected from smartphones only — i.e., a richer set of sensor data from wearable devices can also help (e.g., physiological signal from a smart watch, a vision sensor from a smart glass). We expect data fusion will become more important as a number of commodity wearable devices has been appearing and getting highlights recently. For example, if ECG signals can be reliably captured, it can significantly augment smartphone sensors to infer stress and/or sleep quality for a person [6]. To fully utilise such sensors, the PA engine first needs to provide a unified sensor broker for heterogeneous wearable sensors. The state-of-the art smartphones already have capability to connect to a variety of external sensors over wired and wireless channels. However, ensuring proper interaction with diverse devices can be burdensome, especially when individual application developers need to integrate with multiple sensing devices using various communication medium and data formats. The PA engine should support device discovery, communication, and data buffers so that diverse sensors can be easily integrated to our engine.

4.2 Accurate Inference of HCC

A key requirement for the PA engine is the ability to accurately infer a useful set of human-centric contexts such as interest, preference, intention, emotion, and stress from low-level primitive inputs (such as activity, online posts, and call records) that are easily collectable from commodity mobile and wearable devices. However, inferring human-centric contexts from these lower level inputs is still very much research in progress. In this section, we highlight some of the key research challenges that need to be overcome in this area.

Combining multiple sensor sources and modalities. For the PA layer to be effective, it must combine inputs from multiple sources such as accelerometers, online posts, phone logs, location traces etc. However, how do we combine these inputs (which have different latencies and accuracies)?

For example, prior work has shown that stress can be detected from vocal tones, which could be sensed and obtained through smartphone-embedded microphone [10]. However, for accurate and available inference of stress in diverse life-situations (even in a noisy cafe or restaurant), we should intelligently combine diverse sensing modalities and adjust the weight of different features computed. For instance, we can leverage smartphone-integrated ECG sensors [6] to infer the stress during daily phone usage. We can also use a galvanic skin resistance (GSR) sensor embedded in a smart watch to detect stress levels by analysing skin conductance. Finally, we can also infer stress by observing the user's posts on social networks. However, how do we combine all these various sources together and what weights do we assign to them? This type of multi-modal sensing capability is necessary for the PA engine to be useful. Our current approach is to work with multi-disciplinary researchers (from psychology, organisational behaviour, etc.) to understand which inputs are likely to predict various HCCs and then test and refine our prediction mechanisms through both in-lab user studies and real-world tests using testbeds such as LiveLabs [2].

Error propagation and heterogeneity. A key challenge for the PA layer is handling cascading errors and heterogeneity. For example, to detect stress, the PA layer could be using inputs from the accelerometer, location traces, social network posts, call logs, microphone recordings, and GSR sensor. However, all these various sensors have different error probabilities and distributions. For example, the location system has an error of $\pm$ 5 metres while the social network post sensor has a latency error of up to 2 minutes and a text understanding error of 10%.

A key research challenge will be developing mechanisms to combine these different sensing errors together so that a) the final output is still useful, and b) the error of the final output is understandable. For example, if the PA engine outputs that the user is currently stressed, what is the error in that prediction (if the prediction used 6 different sensors — each with different errors?). Without these error bounds, applications might not trust the output of the engine.

Compounding this problem is the realisation that individual users are very different and they carry a huge number of different phones and sensing devices. For example, the accelerometer sensor in a Galaxy S III has a very different error characteristic from the same sensor in an iPhone 5S. How should the system handle this heterogeneity? A naive solution would be to require per-user per-device training. However, this will not scale. Another key research question would thus be building the techniques that will allow the PA engine to work across multiple users, devices, and sensors with minimal changes and training effort required.

Using contextual information to infer correctness. A solution to the error propagation problem may be using contextual information to identify which sensor streams are most accurate at any point in time. For example, a user's heart rate will be high right after exercise but not when they have been sitting for an hour. Thus, the PA engine should not label a user as nervous or impatient after exercising just by looking at the heart rate (although that might suffice when the user is sitting). Thus, another research challenge is to build an intelligent interpretation layer, that can use a broader set of contextual information (such as location, activity, etc.) to determine the sensing modalities that are most likely to correctly predict the user's current HCC. Again, the challenge will be to solve this problem without requiring extensive user or environment specific customisations.

Energy-efficient data processing. Finally, the PA engine must generate the required HCCs in an energy-efficient manner even when multiple concurrent applications that use different sensing sources and fidelities are running. For example, we can optimise the energy consumption of the *CanICall* application. However, if *CanICall*, and *WhoToHelp* are running concurrently, the energy saving techniques for *CanICall* might interfere with *WhoToHelp* and vice versa. We believe that optimising the energy consumption of multiple sensing applications while preserving all require accuracy and latency levels is still an open problem.

4.3 Long-term Personal Analytics

Many applications will need longer term history HCC trends for the users they are supporting – for example, the *StressReliever* application to trigger 'stress alarms" needs to differentiate acute stress caused by sudden nearby incidents from longer-term stress that are accumulated by uncomfortable human relationships or tight work deadline. This distinction can only by discovered by observing multiple HCCs in a combined way over multiple days and even weeks. To these types of longer-term HCC analysis, the PA engine must 1) develop scalable and energy-efficient data storage techniques that are amenable to various types of historical and long-term analysis, 2) support sufficient privacy for this data and associated queries, and 3) provide a suite of appropriate analysis tools (such as long-term cause-effect analysis).

Scalable and energy-efficient data storage. The first challenge is to store voluminous, continuously streaming sensory readings into appropriate storage repositories (e.g., local, cloud, cloudlets, etc.). Note that local storage will be essential for privacy-sensitive users, and even with cloud use, it is also important to minimise battery and monetary costs to transmit over 3G/4G networks in real time. The "obvious" solution of using in-phone databases (e.g. SQLite) does not work well as these databases are optimised for structured read access while sensing applications frequently have write-plenty read-rarely unstructured workloads. Thus, we need to develop 1) data structure that are suitable to store large amounts of fairly unstructured sensor data, 2) efficient data store operations with proper indexing to enable fast data access, and 3) data ageing and abstraction algorithms that reduces the storage size (and increases privacy) while preserving the capability to answer future queries.

Privacy-awareness in providing HCCs. Letting third-party application query the HCC history and access raw sensor data inherently raises privacy concerns as malicious applications could misuse these functionalities (e.g., reporting work-related depression to management). One solution might be to allow access only to trusted applications or to explicitly request user permissions whenever necessary. However, in reality, these solutions break down when the number of applications and requests increase (and users stop paying attention as well). A more compelling research agenda would be to automatically infer the appropriate privacy rules for specific HCCs based on the context. However, this is very hard even though various studies, such as Klasnja et. al [7] and Choe et. al [3], have shown that context is vital for figuring out the appropriate privacy controls. Another interesting direction is to modify real-time information tracing facilities, such as TaintDroid [5], to monitor unexpected usage of private HCCs.

Long-term cause-effect analysis. A key requirement for the PA engine is to provide analytics support for long-term HCC pattern identification. This type of data analytics support is a very active area of research at the moment – especially for corporate (finance, static customer profiles etc.) or scientific (weather, animal movements, etc.) data sets. However, the key difference with the PA layer is that our analytics need to be performed on mobile data which is a) real-time, b) multi-attribute, and c) from a combination of online, sensors, and phone sources. However, discovering longer-term correlations can be very useful to applications such as detecting the root causes of stress, etc. We are currently building a set of analytics tools that are driven by the concrete problem of stress and obesity among the student population.

5. CURRENT STATUS

We have been building various part of the PA layer and deploying and testing these components using our LiveLabs mobile testbed [2]. Our current status is summarised below. Note: much of this work is still in very early stages and a large amount of research effect still needs to be done.

Data collection: We have been developing an energy-efficient data collector that collects various sensor data (inertial sensors, barometer, sound, WiFi scans etc.) as well as phone usage activities (call/SMS records, application use history, web browsing history, etc.). Our collector supports both Android and iOS devices and is designed to support multiple concurrent applications that require the collection of various combinations of data. We have actively studied, tried (and will continue to do so) various energy optimisation techniques to minimise the energy cost of the collector to a very reasonable level ($<5\%$ amortised energy cost of the total battery power). To date, 800+ undergraduate students at Singapore Management University have provided IRB-approved consent, and downloaded and used our data collector; with about 200+ users uploading their data daily. To collect richer datasets, we are planning to integrate online data sources (Twitter, Facebook activity etc.) with our phone-derived data sources. Finally, we are starting to study the willingness of users to allow higher energy data collection in order to achieve better application performance. This will allow us to understand the situations where a user would be okay with providing more accurate HCCs at the cost of higher energy consumption.

HCC inference: We have taken initial steps to make HCC inferencing possible. First, we have started trying to detect one specific HCC, "buying intention" in a local food court scenario. We first inferred a user's micro movements from low-level sensor data (such as accelerometer and compass), and use that to infer if that user was just passing by, or actively interested in various food items in the food court (and if so, which types of food). Our initial results is that we are to differentiate the food buying intentions of some types of users (vegetarians for example) with a fairly high level of accuracy.

Second, we are actively trying to detect if the user is in a group, and if so, who specifically they are with. This type of group status is a very useful HCC as who you are with (friends, family, etc.) greatly affects your current behaviour (where you go, what you buy, etc.). Our current research attempts to detect groups even without the availability of any form of location data — instead we rely solely on low-level relative sensors such as gyroscopes, compass etc. Our initial results are quite promising and we can identify groups, in some situations, with more than 90% accuracy in under 10 minutes of observation time (even in large crowded Asian shopping malls). Finally, we have begun testing external sensor devices, e.g., ECG sensors integrated in the phone [6], that are more unobtrusive to mobile users. Using these types of direct physiological signals, we hope to augment smartphone-only techniques to improve detection accuracy and coverage.

Long-term personal analytics: To date, we have collected more than 6 month of data from LiveLabs participants. We have started analysing and visualising this dataset to discover interesting behavioural patterns such as how many places does a user visit per day, where are the significant and popular places that a particular user visits, how much does a user interact with other people (using our group detection technique), etc. We are also applying diverse anomaly detection algorithms and performing cause-effect analysis to discover the hidden details in our data set that could provide a deeper understanding of user behaviour.

6. CONCLUSION

In this paper, we have motivated the need for a personal analytics layer, presented a proposed software architecture, and described some of the key challenges in building this layer. We believe that this is an important first step in enabling many different types of new and compelling context-aware applications. In particular, this solves a current problem where application developers also have to understanding sensor data reasonably well if they want to build a context-aware application. By partitioning the responsibilities of sensing the user state from the using of that state for some goal, we hope to make it much easier to build context-aware mobile applications that are both accurate and interesting. We have already started building key components of the PA layer (as described above) and we hope to provide more details of our success (and failures) in the upcoming months.

7. REFERENCES

[1] Balan, R. K., Lee, Y., Wee, T. K., and Misra, A. Challenges in continuous mobile sensing. *Proc. of ComsNets*, Bangalore, India, 2014.

[2] Balan, R. K., Misra, A., and Lee, Y. LiveLabs: Building an in-situ real-time mobile experimentation testbed. *Proc. of HotMobile*, Santa Barbara, CA, 2014.

[3] Choe, E. K., Consolvo, S., Jung, J., Harrison, B., and Kientz, J. Living in a glass house: A survey of private moments in the home. *Proc. of UbiComp*, Beijing, China, 2011.

[4] Chu, D., Lane, N. D., Lai, T. T.-T., Pang, C., Meng, X., Guo, Q., Li, F., and Zhao, F. Balancing energy, latency and accuracy for mobile sensor data classification. *Proc. of SenSys*, Seattle, WA, 2011.

[5] Enck, W., Gilbert, P., Chun, B., Cox, L. P., Jung, J., McDaniel, P., and Sheth, A. N. Taintdroid: An information-flow tracking system for realtime privacy. *Proc. of OSDI*, Vancouver, Canada, 2010.

[6] Kang, S., Kwon, S., Yoo, C., Seo, S., Park, K., Song, J., and Lee, Y. Sinabro: Sinabro: Opportunistic and unobtrusive mobile electrocardiogram monitoring system. *Proc. of HotMobile*, Santa Barbara, CA, 2014.

[7] Klasnja, P., Consolvo, S., Choudhury, T., Beckwith, R., and Hightower, J. Exploring privacy concerns about personal sensing. *Proc. of Pervasive*, Nara, Japan, 2009.

[8] Lee, Y., Iyengar, S. S., Min, C., Ju, Y., Kang, S., Park, T., Lee, J., Rhee, Y., and Song, J. Mobicon: A mobile context-monitoring platform. *Commun. ACM*, 55(3):54–65, Mar. 2012.

[9] LiKamWa, R., Liu, Y., Lane, N. D., and Zhong, L. Moodscope: Building a mood sensor from smartphone usage patterns. *Proc. of MobiSys*, Taipei, Taiwan, 2013.

[10] Lu, H., Frauendorfer, D., Rabbi, M., Mast, M. S., Chittaranjan, G. T., Campbell, A. T., Gatica-Perez, D., and Choudhury, T. Stresssense: Detecting stress in unconstrained acoustic environments using smartphones. *Proc. of UbiComp*, Pittsburgh, PA, 2012.

[11] Nath, S. Ace: Exploiting correlation for energy-efficient and continuous context sensing. *Proc. of MobiSys*, Low Wood Bay, UK, 2012.

[12] Rachuri, K. K., Musolesi, M., Mascolo, C., Rentfrow, P. J., Longworth, C., and Aucinas, A. Emotionsense: A mobile phones based adaptive platform for experimental social psychology research. *Proc. of UbiComp*, Copenhagen, Denmark, 2010.