What am I Doing Now? Pythia: A Mobile Service for Spatial Behavior Analysis

Amnon Dekel, Tomer Weller, Hanny Bar, Cadan Ojalvo
Department of Software Engineering
Shenkar: Engineering, Design, Art
amnoid@gmail.com, tomer.weller@gmail.com, barhanny@gmail.com, cadan85@gmail.com

Scott Kirkpatrick, Benjamin Kessler
Selim Benin School of Engineering and Computer Science
The Hebrew University, Jerusalem, Israel
kirk@cs.huji.ac.il, benjy.kessler@gmail.com

***Abstract*—Pythia is a prototype hybrid Mobile/Cloud service for ascertaining what the user of a mobile phone is currently doing. The service continuously captures and uploads context data to an analysis engine in the cloud. Early field-testing showed that the service can categorize activities into working, at home, traveling, or shopping.**

Keywords-Mobile Context Awareness; Mobile Context Capture; Mobile Applications; Mobile Services; Cloud Services.

I. INTRODUCTION

In the last few years, the use of context awareness as a way of enabling a mobile application to react and perform in relevant ways to the needs and expectations of users has picked up [12][14]. With mobile smart phones containing multiple sensors, the platform has become more capable of capturing data that can be used to identify context. A number of applications have emerged in the last few years that attempted to use this data in order to release the user from having to specifically tell the application what they are doing [2][5][11][13]. Being able to discover the user context is valuable and can enhance and streamline the service being offered.

The work in the area has generated a number of specialized mobile applications on the major smart phone platforms (Android, IOS, Windows phone) that claim to capture signals and recognize user context in an effort to streamline the user experience within specific use cases. The major use cases that have been implemented are driving versus parking in the physical domain (and in the process remembering where the user last parked their car) [2], the personalization of a user's home screen depending on the currently understood context (i.e., presenting a specific set of application icons on the phone desktop when at work versus a different set of application icons when at home [5]), and the presentation of search results that are relevant for the user's current context [11].

We tested some of the available services and found them to be an interesting start, but also lacking in consistent context recognition or transparency. For example, the AGENT [2] application suffered from too many false negatives when parking, leaving the phone in a driving context after we had left the car and thus was not able to remember where the car was last parked. The personalized desktop application COVER [5] that presents a differing set of home screen application icons that it deems to be relevant to the current user context, caused us to feel confused about where the needed applications icons were located on the home screen. This was probably the result of a drastic change in context without the user being an active participant in the change.

Two current services which provide a more consistent and clear experience are MOVES [13] and Google Now [11]. The MOVES application is an activity tracking application that tracks it's users "everyday life" [13]. By gathering sensor based data, it can track a user's activities such as walking, cycling and running, while identifying where these activities take place. It then visualizes its user's day. Google NOW is the most ubiquitous context based mobile service today since it is part of the ubiquitous Android operating system. It is a sophisticated perpetual search service on android that continuously analyses a user's searches, email, calendar and location and uses these signals to present contextually relevant information cards at appropriate times. For example, in the morning it presents a card showing how long it will take to get to work in the current traffic situation (see Figure 1).

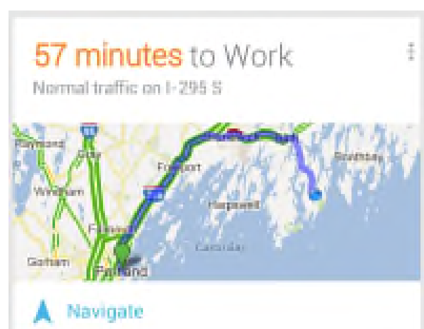

Figure 1. Google Now Context Relevant Card

In both the MOVES application and Google NOW, the actual methodologies for data capture and analysis are proprietary and therefore remain the property of the companies developing them. Although Google makes some of these capabilities available to developers via Application

Programming Interfaces that it opens every once in a while [8], the algorithms remain proprietary, secret and under the control of their owners.

In this paper we present the Pythia system. We start with defining the objectives of the project and then describe the Pythia system. We continue with a description of the methods we used to test the system and end with conclusions and a description of future work to be carried out.

II. Objectives

With Pythia, we are striving to develop a mobile and cloud based service for capturing contextual signals and ascertaining the current user context. Because of a lack of context capture consistency and transparency in existing mobile applications, the goal of the project is to learn about and improve the capability of a mobile-based service to arrive at better context classifications.

Pythia was developed as part of a research effort to tap the intimate relationship that people have with their phones in order to learn about their activities and to use this knowledge in order to make the phone a smarter companion in our everyday lives.

III. The Pythia System

Pythia is a hybrid mobile/cloud service that captures packages and uploads data to a web-based repository where the data is parsed, cleaned, normalized and classified. All Pythia data is stored on a MongoDB instance. The Pythia server is a lightweight Node.js application running in a hosted Heroku service. Data is hosted on MongoHQ (see Figure 2).

The phone client includes a management User Interface (see Figure 3) to enable the user to set up the data capture resolution and upload dynamics (i.e., what is the minimal distance traveled that the system will save and upload to the online service), and a background service (see Figure 4) that gathers location data and uploads it to the data repository according to the application settings. The background service shows the amount of events it has captured and is ready to upload to the server on the next synchronization process ("events in pipe").

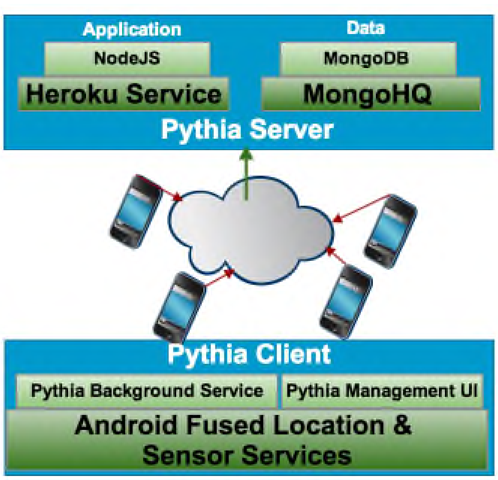

Figure 2. Pythia System Overview

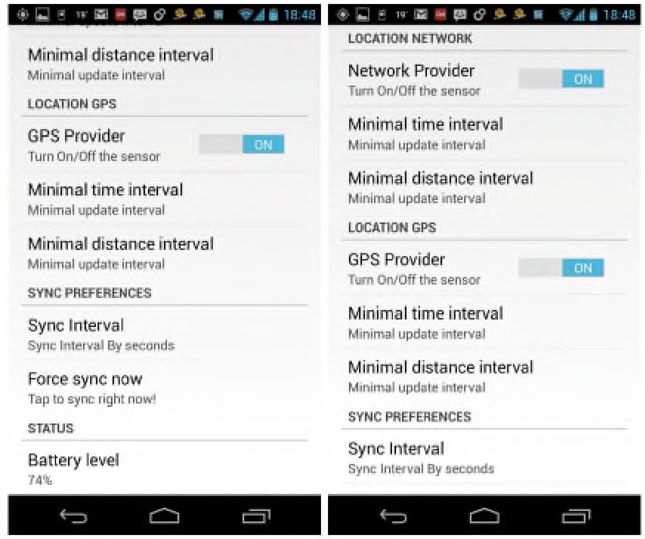

Figure 3. Pythia Management UI

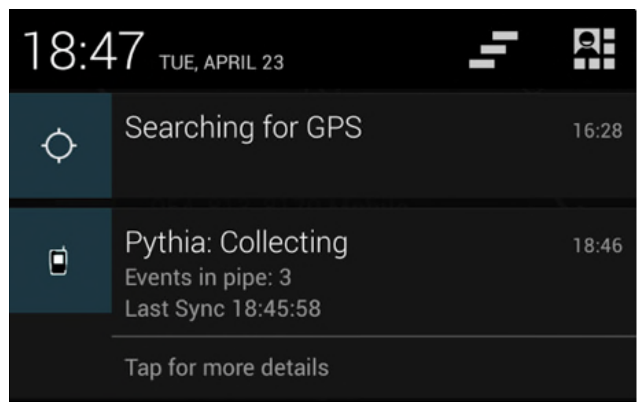

Figure 4. Pythia background Service

Note that the management UI was designed for our initial technical field trial and is not designed for normal users. Future iterations of the service will be released to larger numbers of users and will include a more streamlined and simple user interface that will be easy to use by anyone.

The Pythia server receives the sensor data (GPS or Cell based location), from the participating mobile clients and performs a spatial classification process on the data. This is achieved with a series of three map/reduce [6] computations:

1. **Pre-processing:** Similar location events are grouped by their bounding timestamp and geobox and then these aggregated events are indexed using geohashing- a method of transforming a 2d point into a 1d hash, which allows for easy indexing and simple geoboxing of locations. This process reduces the size of the data while filtering out noise caused due to the instability of the mobile device location reads.

2. **"HotSpot" aggregation:** Grouping of aggregated location events by bounding geobox, and

 ISBN: 978-1-61208-366-7

calculating the sum of timeframes per geobox and then sorting by this sum. Geoboxes with highest timeframe count represent the users most visited locations.

3. **"HotSpot-by-hour-slot" aggregation:** Same as "HotSpots" but also assigning time of day to geoboxes. Allows us to see the user's most visited locations per hour.

Once the sensor data per phone was processed as described, the next step was to ascertain a context for each location and timeslot. For this version, we kept the context to the following: *Home*, *Travel*, *Work* and *Shopping*. It is clear that this classification is very rough, but we think that being able to identify these prototypical contexts is a valuable step in the process of being able to identify more refined contexts later on.

IV. METHOD

An initial version of the Pythia data capture service was installed on two Motorola Razr's and two Nexus 4 phones. The devices continued to be used as active phones for a week while capturing and uploading the data to the Pythia server.

In this version, the service was only used to gather location and sensor data. Apart from being able to turn the service on or off, or to tweak the sensor gathering resolution, the user did not get any information from the application itself. The uploaded data was then processed at the server end with the goal of ascertaining what the user was doing throughout the period of data capture.

The goal of the service was to establish if it was feasible to use the sensor information in order to conclude what context the user is in, while keeping within a viable power consumption envelope. It is clear to us that a service that can conclude a context correctly, while using too much power, will not be useful, and at the same time, a service that is very power efficient but cannot ascertain the current context is not useful either. The sweet spot is to be power efficient while correctly ascertaining context. We define power efficiency as a power consumption envelope that does not visibly lower the service life of a smart phone in normal use to below a full day of use. A power consumption envelope that lowers the daily service life of a phone will simply be discarded by most users.

In this version, the process of ascertaining user activity context was semi-automatic and necessitated human supervision. The service would automatically aggregate and run through the list of locations and timeslots and then receive human supervision as to the naming of the current context. Once enough of these were seen by the system it was able to classify contexts. Thus, after seeing that a specific location was always classified as Home by the supervisor, the service would be able to reach the same conclusion going forward.

V. INITIAL RESULTS

The first version of the service was quickly seen to be overly power hungry and severely shortened the daily service life of the phone by 50%. This was unacceptable. Google then released a new sensor capture framework (called Fused Location Provider [10]) that was purported to be more power efficient. After implementing it in the Pythia client we were able to lengthen the service time to the full day threshold we expected and enabled us to continue into a 4-week field test with the 4 phones that continued to be used in a normal fashion by their owners.

Figure 5 shows the distribution of classified activity contexts over a representative 24-hour period. For each of the participants, we compared the resulting activity map to their calendar and notes as well as interviews. The comparison showed us that the classifications were correct.

Figure 6 shows the total averaged activity distribution as identified over the course of the process. In this case, we averaged identified activities over the time of the trial and the most frequent activity per timeslot became the selected classification. These were mapped onto a 24-hour representative period. Note that we analyzed only normal working days, ignoring weekends and vacations.

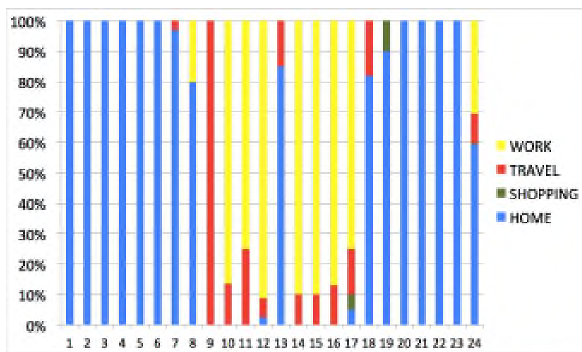

Figure 5. Classified Activity Context over 24 hours
Y-axis: the percentage of each activity during a specific hour
X-axis: The hour in the day

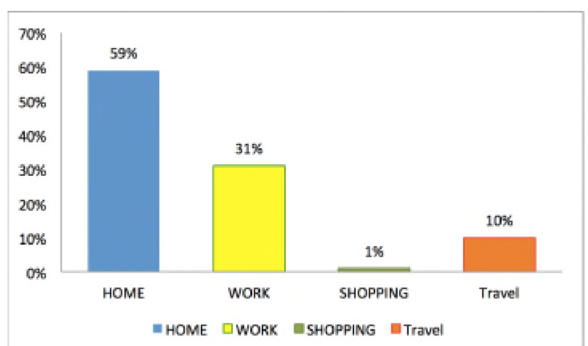

Figure 6. Total Activity Distribution

Analyzing the results presented above, we conclude that within the constraints described, the system was able to classify the major activities across time.

VI. Conclusion And Future Work

The Pythia system analyzed an aggregated database of user locations over time and used this in order to classify what the user was doing at each location during each time slot. We believe that being able to semi-automatically identify what a user is doing (their personal context) is an important signal that can help in making our mobile phones more useful and helpful in our daily lives.

While we have shown that the Pythia system is able to classify activity contexts related to location and time, we believe that ascertaining context via location alone may be too coarse as a signifier in many cases. Because of this, we are now in the processes of improving the system in the following ways:

1. Widening the net of possible activity context classifications that the service can identify.
2. Minimizing the need for human supervision in the classification.

We believe that both of these improvements can be achieved by supporting additional sensors in the system. Additional *physical sensors* can help to an extent, i.e., using physical activity recognition to identify if the user is stationary, walking, running, riding a bike or travelling in a vehicle, is valuable, but has a limited capability for refined context discovery. Similarly, the use of audio analysis can help us identify how many people are around the user, if they are in a meeting or in a social or sporting event. This can also be achieved by identifying the wireless fingerprints of radio devices around a person (wireless hotspots, mobile phones, etc.). But we believe that fusing sensors of different types will create the most value. For example, by adding *semantic sensors* such as the phone calendar or various text-messaging services, we gain both a wider classification capability, and also a method that can minimize the human supervision needed. When a person adds a calendar entry (for example "Meeting with Joe" or "Pick up Anna from school"), the descriptive text can serve as a semantic data point that helps identify a more refined context (i.e., “Work Meeting” or “Family Task”) while at the same time serving as an automatic supervision entry. Similarly, a calendar entry named “workout” combined with physical sensors that identify the user as in a running activity out of doors will help in identifying the context as “Jogging”.

An updated version that we are now working on (Pythia Occursum) will include the calendar as our first semantic sensor. Additional sematic sensors that will be added are the contacts database, the text-messaging database(s), and the user's email.

Acknowledgment

This research was funded in part by the Intel Collaborative Research Institute for Computational Intelligence (ICRI-CI).

References

[1] A. Abecker, A. Bernardi, K. Hinkelmann, O. Kühn, and M. Sintek, “Context-aware, proactive delivery of task-specific information: The KnowMore project,” Information Systems Frontiers, 2(3), 2000, pp. 253–276.

[2] “Agent Mobile Application,” On the Google Play store: [https://play.google.com/store/apps/details?id=com.tryagent]

[3] P. Bellavista, A. Corradi, M. Fanelli, and L. Foschini, “A survey of context data distribution for mobile ubiquitous systems,” ACM Comput. Surv. 44, 4, Article 24, August 2012.

[4] C. Bettini, O. Brdiczka, K. Henricksen, J. Indulska, D. Nicklas, A. Ranganathan, and D. Riboni, “A survey of context modelling and reasoning techniques,” Pervasive and Mobile Computing 6 (2), 2010, pp. 161–180.

[5] “Cover Mobile Application,” On the Google Play store: [https://play.google.com/store/apps/details?id=com.coverscreen.cover]

[6] D. Jeffrey and S. Ghemawat, "MapReduce: simplified data processing on large clusters," Communications of the ACM 51.1, 2008, pp. 107-113.

[7] AK. Dey and GD. Abowd, “Towards a better understanding of context and context-awareness,” CHI'2000 Workshop on the What, Who, Where, When, and How of Context-Awareness, 2000, pp. 304-307 [https://smartech.gatech.edu/ bitstream/ handle/ 1853/ 3389/ 99-22.pdf]

[8] A. Duvander, “Google Now for Android Full of APIs,” On the Web at [http://blog.programmableweb.com/2012/06/27/ google-now-for-android-full-of-apis/]

[9] P. Fahy and C. Siobhan, "CASS–a middleware for mobile context-aware applications," Workshop on Context Awareness, MobiSys, 2004.

[10] “Google Fused Location Provider Application Programming Interface,” on the Android Developer web site: [http://developer.android.com/google/play-services/ location.html]

[11] Google Inc, “Google Now,” On the web at [http://www.google.co.il/landing/now/]

[12] R. Kyle, “The Future Of Mobile Is Figuring Out What You Want Before You Do,” On the Web at Business Insider: [http://www.businessinsider.com/the-future-of-mobile-is-context-2014-2?op=1]

[13] “Moves Mobile Application,” On the Google Play store : [https://play.google.com/store/apps/details?id=com.protogeo. moves]

[14] S. Perez, “T3 Reveals Scout, A Mobile “Context Engine” That Knows Where You Are & What You’re Doing, To Personalize Apps & Sites,” On the web at Techcrunch [http://techcrunch.com/2013/08/06/t3-reveals-scout-a-mobile-context-engine-that-knows-where-you-are-what-youre-doing-to-personalize-apps-sites/]

[15] A. Schmidt, M. Beigl, and H. Gellersen, "There is more to context than location," Computers & Graphics 23.6, 1999, pp. 893-901.

[16] Y. Tingxin, D. Chu, D. Ganesan, A. Kansal, and J. Liu, “Fast App Launching for Mobile Devices Using Predictive User Context,” MobiSys 2012, Low Wood Bay, UK, Jun 25-29, 2012, pp. 113-126.

[17] K. Wan, “A Brief History of Context,” IJCSI International Journal of Computer Science Issues, Vol. 6, No. 2, 2009, pp. 33-42