Context Recognition on Mobiles*

Fatih Tümen
DIVA Group
University of Fribourg
CH-1700 Fribourg, Switzerland
fatih.tuemen@unifr.ch

ABSTRACT

Increasing importance of context awareness in ubiquitous computing put mobile devices on target as they have the most tempting infrastructure for context recognition. This report presents state of the art in context recognition on mobile devices with overview of some of the most popular applications in this domain. We present discussion of most prominent aspects such as hardware availability, ethics,o pen problems, architectural and design challenges in developing context recognizing applications on mobile devices.

Keywords

Multimodal interaction, context recognition

1. INTRODUCTION

Since 90's context awareness increasingly gained importance in ubiquitous computing as it seems to be a promising solutions for a lot of problems in the ever-changing environments of nowadays [8]. Simply put, context awareness is being smart which requires adaptation to the environment. This adaptation entails on the answers of questions such as who, what, when, where and how. Context recognition is the process of finding these answers by deducting interpretation, understanding, and abstraction of context, using either or combination of data acquired through sensory systems, intermediate information, or even the context itself previously perceived.

Until about a decade ago the word context had been used as complementary to the location. Modern realization of context, however, has a broader definition [1]. Context can be any kind of information including but not limited to location, time, activity, identity or any other entity that can help characterize parameters of a system or make inferences in order to enhance system's behavior thereby the user experience. Such broad understanding of context makes it very much application dependent. A context in one system can be an input data to derive context from in another system or not relevant at all. Therefore categorization of context with respect to system's needs plays a key role in design of context aware systems.

Recent advancements in mobile computing, increasing number of sensors on mobile devices, application programmable interfaces to these sensors, open source mobile operating systems and libraries have altogether provided a tempting platform for developing context aware applications. While these advancements offer diverse set of sensory data to use for context recognition, they come with limitations and brought various challenges to tackle [5].

In this paper we will present an overview of the state of the art in context recognition on mobile devices. We will study some of most popular context aware applications and discuss their use of current mobile technologies, their approach to tackle encountered challenges and open problems in context recognition. We will begin with discussion of design principles one should consider while developing context recognizing applications in Sect 2. and follow up with discussion of preferred context recognition architectures in Sect 3. Discussion of open problems in the field and of future work can be found in conclusions section.

2. DESIGN PRINCIPLES

While mobile devices continue to provide more computation, memory, storage, sensing, and communication bandwidth, they are still resource limited for context recognition process which require relatively complex signal processing and machine learning algorithms [5]. Selection of target platform is correlated to target user as well as sensors needed.

The target platform, thereby their resource constraints can solely determine the algorithms to be used in context recognition. Most mobile devices provide application programmable interfaces to their hardware and software development kit to develop context aware applications. However, different hardware and APIs are offered by different manufacturers [5]. This diversity introduces various challenges to tackle.

For developers targeting multiple platforms this makes it difficult to port their application across platforms. ContextPhone [7] addresses this very problem by providing hardware

*This report is produced for Seminar on Multimodal Interaction on Mobiles Devices https://diuf.unifr.ch/main/diva/teaching/seminars/seminar-multimodal-interaction -mobiles-devices abstraction for a higher level applications needing a specific functionality of a sensor.

2.1 Hardware

Sensors in mobile devices are introduced mainly to enhance user experience. After the introduction of application programmable interfaces to these sensors it has become common to use them, usually in combination with each other, for context recognition [5].

Some of the most commonly used sensors in context recognition include accelerometers, proximity sensors, light sensors, Global Positioning System (GPS), compass, gyroscope, microphone, camera, and Bluetooth. It can be tempting to use as much data from sensors as possible for high quality context recognition, But this comes with a cost of higher power consumption. The more sensors an application uses the shorter the batter lasts.

Finding an energy efficient algorithms for using sensors is still an open problem of today. Therefore developers should begin with determining possible context they are interested in and the mere sensors that can be used to acquire data from which these context can be derived. More detailed discussion of sensors will follow in Sect 3.

2.2 Ethics

Context can be any kind of information but can we collect any information we like? Privacy policies can determine whether the system will be used or not. The users may simply choose not to if they feel their privacy is compromised. Despite the available cryptographic tools to preserve user identity, people remain sensitive about how personal data is collected and used especially.

A common approach in systems with online backend is to upload not raw sensor data but extracted features only. Processing sensory data entirely on the mobile may be less appealing in terms of limited resources but can offer privacy advantages [6].

3. ARCHITECTURE

Context recognition on mobile devices is rather a new domain and it is no surprise to see that there is no conventional architecture [5]. Main reason for this is diversity in hardware and its consequences mentioned in the preceding section.

In order to give a context to remainder of this section, this is a good place to define general idea behind context recognition. In essence context recognition happens in three main stages. First, the data is retrieved from sensors. Number of sensors used defines input modality. The sensor data is then used to extract features. Feature extraction process can be based on simple statistics as well as complex mappings and algorithms. Lastly these features extracted are used to derive context. Separation of these steps is key to a robust architecture. Such abstraction can make it possible to deploy individual modules to handle different stages [4].

3.1 Modality

It has been argued that most successful inferences can be made from multi-modal sensor streams [5]. It is difficult for us to make a comparison in this aspect because the only single modal application we have been discussing is SoundSense [6] which presents successful context recognition results. Lu et al. in their work SoundSense uses single sensor data from microphone to classify sound events. Their success lays in bringing the user into hybrid supervised and unsupervised learning cycle which they call adaptive unsupervised learning algorithm. However, this approach relying on intermediate user prompt has a negative effects to the self-autonomy of the system. A counter solution is prompting other systems rather than the users. We will further elaborate on this later in this section.

3.2 Sensors

Today's mobile devices offer wide range of sensors and interfaces to these sensors for sensory data acquisition. Using more sensory data gives context recognition a higher accuracy at the price of higher power consumption.

Context can be derived from single or multiple sensory data. While using single sensor can be enough to make an inference successfully [6], bringing more sensors into action is also a common practice [2, 3, 4]. Let us list commonly used sensors and their use cases in mobile devices for context recognition:

- *Accelerometers* are used usually to rotate the screen with respect to the orientation in which user is holding the device.
- *Proximity sensors* are used to estimate how close the user is to the device. In combination with accelerometers they can be used to change device mode, for instance, from SMS reading to phone call mode.
- *Light sensors* can sense the level of ambient lighting and help adjust the brightness of the screen accordingly.
- *GPS (Global Positioning System) sensors* allow systems to localize itself in world. They can be as accurate as 2-3 meters in distance and 14 nanoseconds in time. Continuous GPS data can tell the difference, for example, between walking or driving by car and help recognize activity context.
- *Compass and gyroscope* provide directional and rotational information in addition to positional information provided by GPS sensor. All these three can be combined to improve quality of location context and produce geographically relevant output.
- *Microphone* is probably the least exploited yet one of the oldest sensors in mobile devices. Lu et al. [6] in their application SoundSense use this sensor to recognize context by listening and classifying ambient sounds continuously.
- *Camera* in the rear of mobile devices is traditionally used to have a memory of the picturesque. The front camera, now commonly found in devices supporting 3G network, where it is used for video conferencing, can also be used to recognize faces for authentication or to track eye movements and turn the page in e-book reading mode or simply to detect presence and attention of the user.

- *Bluetooth* is a short-range frequency wireless protocol, that has been primarily designed for cordless headset accessories, can also provide information about other bluetooth enabled devices nearby. This latter aspect of bluetooth and long-range RF GSM network are being exploited by Eagle et al [2] to infer activity and location context and recognize patterns in social environments.

There are sensors (e.g. infrared sensor, temperature sensor, touch sensor) which have not yet found a common use in context recognition. There are also sensor which we have not seen on mobiles devices yet, such as humidity sensor or barometer. It has been suggested that [5] barometers, for example, can be used in along with accelerometers to recognize a particular movement such climbing stairs.

One potential use case for this would be in a medical monitoring application that have been increasingly gaining attention on mobile devices. Using medical monitoring applications doctors can monitor their patients' heart rate, blood pressure etc.' remotely. Sensors on the mobile device can be use to bring immediate context to the this data and tell doctors what kind of activity the patient was pursuing before and while the patient had, for instance, a seizure or an increase in blood pressure.

Sensor usage should be handled with care, especially in such real time applications, as it can have dramatic effects on power consumption given that there is limited support provided by vendors on sensor control and power management. Most applications also need multitasking and background processing which is currently supported only by Android and Nokia Maemo phones leaving popular iPhone 4 OS, that support only the notion of multitasking, inadequate.

There is an apparent need for less power greedy algorithms as well as more openness in hardware access. Based on these points it is for developers to investigate available programmable interfaces to sensors found on which platforms and which of these sensors can be used to recognize which context. Mainly the questions (e.g. who, when, where and what) to which system seeks answers to constitutes the categories of context in which the system is interested in (e.g. identity, time, location, activity).

3.3 Events

The events are mappings of contextual information to real life situations using information theory and probability theory to infer context from sensory data. For derivation of contextual information, feature extraction may require more processing in single-modal architecture (e.g. using audio or vision) than in multi-modal architectures. In multi-modal architecture each sensory data contribute a part to the whole picture [4]. Table 1 list how situations and data from different sensors may relate.

Individual systems tend to collect continuous data because it improves context quality in return. Polling sensors continuous requires multitasking and more power. However not all mobile platforms support multitasking. Also due to limited hardware access on current mobile devices, applications cannot perform power management. So they face to polling rate versus battery life trade-off [2]. Lane et al. argue that [5] for continuous sensing to be viable there need to be breakthroughs in low-energy algorithms that duty cycle the device while maintaining the necessary application fidelity.

ContextPhone [7] is a software platform built in this direction. It uses a subscribe-publish approach to notify application about contextual events. Application interested in a particular contextual event subscribe to that event and ContextPhone updates application as the subscribed event occurs.

The aim is to provide subscriber applications a robust and reliable interface for rich date acquisitions with minimum control and maintenance eliminating the need for dealing with hardware level details. This in return provides rapid development of applications. It is mentioned that using such an interface Eagle et al. saved significant time in development of their applications measuring social relationships.

3.4 Scalability

Context recognition systems are expected to survive highly dynamic mobile environments. For example while it can be reasonable to use supervised classification algorithms for small scale domains, it impossible to train a such system for larger scale domains. Thence, for a successful system, scale of the application domain should to determined beforehand.

Different types of data is one of the challenges need to be tackled in mobile environments. One of the solutions is to use a hybrid approach [6] that can handle new types of data and context. Lu et al. uses generic models for general sounds in people's lives and for more specific sounds they use unsupervised learning techniques to discover and learn new sound types with the help of user interaction.

A counter solution is to let mobile devices interact with each other to build new sensing paradigms based on collaborative sensing. This would allow devices to inquire other devices about unknown situations rather than the user. Such approach which is in contrary to the one presented by Lu et al. can be used to generate a social knowledge base that can be used across various services over the cloud.

Personal context recognition applications would just inform the individual, while social networking context recognition application may share inferences, context, or activity with friends which is something required to harness the potential of context recognition [5].

Moreover, user expectations of more intelligent output from a ubiquitous system could be achieved with introduction of uniform context modeling in form of which devices could exchange knowledge on the cloud. As a result, experience and expectation mobile users would be shifted and improved worldwide.

4. CONCLUSIONS

About twenty years ago, when Mark Weiser introduced his vision of ubiquitous computation embodied in ParcTab, the focus of context awareness was mainly on location. Less than a decade later Abowd et. al gave context broader cov-

Table 1: Real world situations related to sensor data [4]

Situation	Sensor Data
User sleeps	It is dark, room temperature, silent, type of location is indoors, time is ‘night-time‘, user is horizontal, specific motion pattern, absolute position is stable
User is watching TV	Light level/color is changing, certain audio level (not silent), room temperature, type of location is indoors, user is mainly stationary
User is cycling	Location type is outdoors, user is sitting, specific motion pattern of legs, absolute position is changing.

erage. The infrastructure available back then was perhaps not enough to fully utilize this definition.

Today the main issue with context aware systems is utilization of the infrastructure rather than availability of it. Millions of people are already carrying mobile devices that are capable of more than what personal computers could do a decade ago. Despite that mobile devices are still resource limited devices for context recognition which can be a computationally expensive task. Cloud computing, to whatever degree, can give this an ease. However there are still remaining challenges to be tackled on the mobile end.

One of the main issues in context recognition on mobiles is lack of openness in hardware level access to mobile sensors which makes power management by applications impossible. In this context development of energy efficient algorithms for sensor data retrieval remains to be an open problem. Perhaps one of most underestimated aspect in context recognition is the privacy. In order to produce better quality context, system will want to know as much as possible about the the user. Before doing so data collection, usage and storage policies must be strictly defined and made transparent to the user.

Although diversity of available hardware may count as a pro to some level, variety of APIs on these different hardware makes it challenging for developers to port their applications across platforms. Currently research groups try and develop their own abstraction layer between sensory layer and context recognition. Community is in obvious need for standard in this field. In order to satisfy community needs, a research consortium could lead the way to standardizing sensor application programmable interfaces and making conventional architectures possible in context recognition. Solution to these problems will give a boost to advancements of context recognition on mobiles, exploiting the potential in new application domains and change people’s social life.

5. REFERENCES

[1] G. D. Abowd, A. K. Dey, P. J. Brown, N. Davies, M. E. Smith, and P. Steggles. Towards a better understanding of context and context-awareness. *CHI 2000 workshop on the what who where when and how of contextawareness*, 4(What, Who, Where, When and How of Context-Awareness):1–6, 1999.

[2] N. Eagle and A. Sandy Pentland. Reality mining: sensing complex social systems. *Personal and Ubiquitous Computing*, 10(4):255–268, 2005.

[3] J. Froehlich, M. Y. Chen, S. Consolvo, B. Harrison, and J. A. Landay. MyExperience. *Proceedings of the 5th international conference on Mobile systems, applications and services - MobiSys ’07*, San Juan,:57, 2007.

[4] H. Gellersen, A. Schmidt, and M. Beigl. Multi-Sensor Context-Awareness in Mobile Devices and Smart Artefacts. *Mobile Networks and Applications*, 7(5):341–351, 2002.

[5] N. D. Lane, E. Miluzzo, H. Lu, D. Peebles, T. Choudhury, and A. T. Campbell. A survey of mobile phone sensing. *IEEE Communications Magazine*, 48(9):140–150, 2010.

[6] H. Lu, W. Pan, N. D. Lane, T. Choudhury, and A. T. Campbell. SoundSense. In *Proceedings of the 7th international conference on Mobile systems, applications, and services - Mobisys ’09*, MobiSys ’09, page 165, New York, New York, USA, 2009. ACM New York, NY, USA, ACM Press.

[7] M. Raento, A. Oulasvirta, R. Petit, and H. Toivonen. ContextPhone: a prototyping platform for context-aware mobile applications. *Ieee Pervasive Computing*, 4(2):51–59, 2005.

[8] T. Strang and C. Linnhoff-Popien. A Context Modeling Survey. *Graphical Models*, Workshop o(4):1–8, 2004.