Memo-it: Don't write your diary, Sense it

Karl Aberer
karl.aberer@epfl.ch

Michele Catasta
michele.catasta@epfl.ch

Georgios Christodoulou
georgios.christodoulou@epfl.ch

Ivan Gavrilovic
ivan.gavrilovic@epfl.ch

Filip Hrisafov
filip.hrisafov@epfl.ch

Mathieu Monney
m.monney@epfl.ch

Abdessalam Ouaazki
abdessalam.ouaazki@epfl.ch

Boris Perovic
boris.perovic@epfl.ch

Horia Radu
horia.radu@epfl.ch

Jean-Eudes Ranvier
jean-eudes.ranvier@epfl.ch

Matteo Vasirani
matteo.vasirani@epfl.ch

Zhixian Yan
zhixian.yan@samsung.com

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the Owner/Author. Copyright is held by the owner/author(s).
UbiComp '14 Adjunct, September 13-17, 2014, Seattle, WA, USA
ACM 978-1-4503-3047-3/14/09.
http://dx.doi.org/10.1145/2638728.2638753

Abstract
The profusion of sensors embedded in modern mobile devices collect an increasing amount of information about the activities performed by a user. Leveraging the episodic memory model defined by neuroscientists, Memo-it exploits this information to create a multi-scale structured representation of the user's activities in a semi-automated fashion, while preserving the privacy of the user's data. In addition to building a digital diary of the user, the semantic approach taken by Memo-it is able to answer multi-dimensional queries, and to enable the inter-operability of memories between users.

Author Keywords
Memo-it; activity recognition; automatic diary

ACM Classification Keywords
H.3.3 [Information storage and retrieval]: Miscellaneous; I.5.1 [Pattern recognition]: Models; H.1.2 [Models and principles]: User/Machine Systems.

Introduction
The amount of user-centric information that can be gathered by recent mobile devices is tremendous. Furthermore, the cost in term of battery consumption to acquire such information is getting lower with every new generation of devices, while the APIs offered by

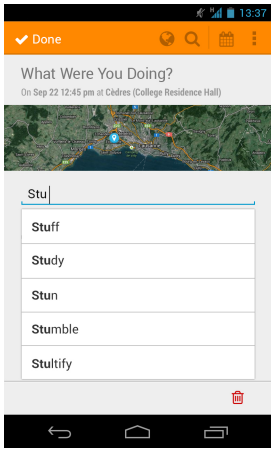

Figure 1: Creation of a new EE. Selection of a verb from the ontology.

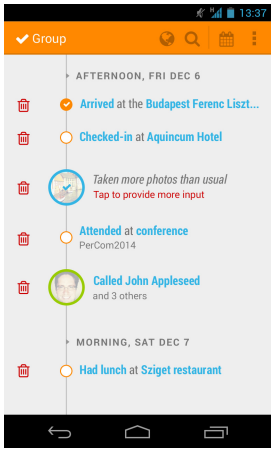

Figure 2: Manual aggregation of the EE.

smartphone OSes make it fairly easy to collect this data.

Memo-it is a mobile application for Android based on the ideas developed in MemorySense [1]. Following the episodic model described by Tulving in 1983 and its organization described by Conway more recently [4], Memo-it performs sensor data fusion and activity recognition to aggregate different sensors readings into digital memories, with minimal intervention of the user.

Related work

Collecting documents and information about a user while offering aggregation and visualization capabilities over such data is a well-explored area. In 1945, at the genesis of the computer era, Vannevar envisioned Memex [3], a hypothetical system in which users would store all the documents related to their lives. In early 2000, Gordon Bell implemented this vision with MyLifeBits [5]. This system would allow not only to store different sources of information, but also to search the database through full text search, annotations and hyperlinks.

More recently, the increase in the number of available data sources is leading to new opportunities for developing more advanced methods [8] and systems [2]. In [7], Rawassizadeh et al. present UbiqLog, a flexible life logging system for mobile platforms. Memo-it improves upon the data model of UbiqLog by exploiting the episodic memory model of Conway, and semantically structuring it. This improvement enables more powerful memory browsing techniques. To the best of our knowledge, Memo-it is the first system that leverages the episodic memory model of the human brain in order to build a semantically enhanced representation of the users' activities.

Mobile sensing framework

The sensing framework of Memo-it is responsible for gathering data from the sensors, and then building the user memories with two primary goals in mind: 1) the application must require as less interaction as possible with the user, while 2) preserving the privacy of the user's data. Indeed, Memo-it aims at building the memories of the user solely based on sensor data, after a learning phase supervised by the user is performed in order to generate the user activity models. Furthermore, Memo-it has full rights on all the private data accessible by the phone. In order to prevent that sensitive data is leaked to third parties, all the computation and all the pre-processed data is kept locally on the device. Therefore, in order to process the data in a mobile environment, the use of energy efficient algorithms is required.

In the episodic memory model defined by Conway, Episodic Elements (EE) represent the atomic bricks of memories. Multiple EEs can then be aggregated into Simple Episodic Memories (SEM), which represent more abstract memories, e.g., the EEs "going to the museum with Alice", "having a coffee with her" and "walking by the lake" can be aggregated into the SEM "going on a date with Alice". Memo-it mimics this model by generating episodic elements from sensor information, and then reconstructing episodic memories from episodic elements.

Generation of episodic elements

Raw sensor data is used to define episodic elements along 3 dimensions *When*? (The time at which the episode happens) *Where*? (at which venue/location it occurs) and *What*? (the episode description).

While *When* is trivially addressed using a timestamp generated by the system, *Where* and *What* needs to be

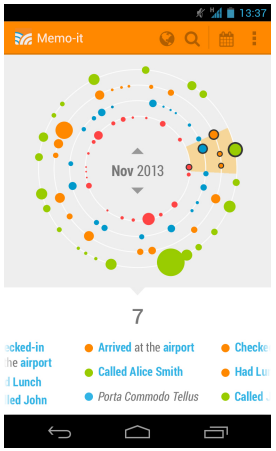

Figure 3: Monthly overview.

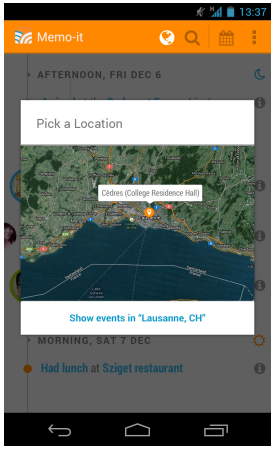

Figure 4: Search based on location.

defined in a structured fashion by the user, or generated by the system automatically.

Where: to define the venue where an episodic element happened, Memo-it detects the points of interest visited by the user. Leveraging on geofencing and the raw GPS coordinates acquired by the phone, Memo-it uses a finite state machine to determine if the user spent more than 10 minutes in the same location.

This technique allows for noise removal by discarding locations where the user just traveled by. In addition, the usage of low accuracy GPS sensing to trigger more accurate GPS sensing is energy efficient (i.e., opportunistic sensing). Once a point of interest is detected, its location is anonymously transmitted to the location-based service Foursquare[1] in order to get a list of nearby venues. A recommender system using item-based filtering determines then which of the venue should be used. The similarity function used for the filtering is a weighted sum of both the number of times the user visited each venue and the distance of each venue to the actual GPS coordinates.

What: The activity description is composed of a verb and a noun taken from the Wordnet database [6]. As shown in Figure 1, the application auto-completes the user input.

Memo-it can infer both the activity description and the location of the event automatically, based on the statistics of the user. However, a training phase is required to build these statistics during which the user has to manually input the event description.

Creation of episodic memories
In order to match the Tulving's episodic memory model, EEs are aggregated into SEMs. Due to the unbounded amount of potential activities and given that the way a complex activity is perceived and executed differs from one user to another, it is not possible to build a knowledge-base which would be exhaustive enough to accommodate every user. Thus, Memo-it follows a 2-phase approach. First a personalized model for each SEM is created based on the user manual input (as displayed in figure 2). Then, when an initial model of a SEM has been generated, Memo-it automatically detects the occurrence of this SEM and, if necessary, updates its model.

For the sake of flexibility and since activities are not always performed in the same way, the models need to allow reordering of the EE and deletion/insertion of a new EE. Another constraint comes from the mobile environment, in which computation power is precious. For these reasons, the model generation needs to be incremental and of low computational complexity.

At the time of writing, two approaches are being evaluated to generate the models. The first approach uses a finite state machine for each SEM, where EE are defined as states and transition probability between states are determined by the encountered instances. The second approach is based on incremental association rules mining (ARM). Since ARM does not support the notion of ordering, two layers of rules are generated: one layer focuses on the presence of specific EEs in the SEM, the other one focuses on the presence of transitions in the SEMs.

User interface

The episodic memories generated by the framework can be offered to the users in different ways. The screenshot displayed in figure 3 shows an overview of a month-worth

[1] http://www.foursquare.com/

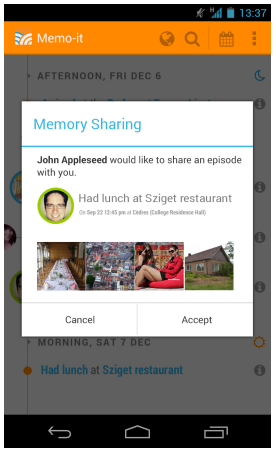

Figure 5: Sharing memories between users.

of user activities. This view allows the user to have, in a glimpse, an idea of when she was active the most (size of the circles) and for what kind of activity (life events, social interactions, phone events).

Search capabilities
The EEs and SEMs of the system are structured using values coming from an emergent ontology. While this approach restraints the degree of freedom of the user when inputting new memories, it enables the system to generate memories in a more structured fashion, as compared to just free-text descriptions. A resulting advantage of this approach is the possibility to efficiently query the system, e.g., retrieving memories that match a certain criteria. It is therefore possible to search along each dimension (e.g., location, activity description, etc.) or a conjunction of dimensions as described in figure 4.

Shared memories
A second feature of Memo-it is the possibility to share memories among users. As displayed in figure 5, by leveraging the low energy Bluetooth technology of recent smartphones, users are able to share memories easily with other users participating to the same event. Sharing memories reduces the number of interactions between the user and the application: in case the application did not generate the memory automatically, only one of the two users needs to input the details of the memory.

Conclusion

Memo-it is a digital diary that can learn from the user's regular pattern to automatically detect and generate her everyday memories. A privacy preserving approach in which all the computation is done on the mobile phone ensures that no sensitive information is leaked to third parties. Finally, the semantic approach used in the application helps the user to browse more efficiently her digital memories.

Acknowledgements

Memo-it is part of the MemorySense project, sponsored by Samsung Research America, and by Haslerstiftung in the context of the Smart World 11005 (MEM0R1ES) project. We are also thankful to Andrea Peltrin for the great design work.

References

[1] Aberer, K., Catasta, M., Radu, H., Ranvier, J.-E., Vasirani, M., and Yan, Z. Memorysense: Reconstructing and ranking user memories on mobile devices. In *Percom WIP '14* (2014), 195–198.

[2] Bicocchi, N., Castelli, G., Mamei, M., Rosi, A., and Zambonelli, F. Supporting location-aware services for mobile users with the whereabouts diary. In *MOBILWARE* (2008), 6.

[3] Bush, V., et al. As we may think. *The atlantic monthly* (1945), 101–108.

[4] Conway, M. A. Episodic memories. *Neuropsychologia* (2009), 2305–2313.

[5] Gemmell, J., Bell, G., Lueder, R., Drucker, S., and Wong, C. Mylifebits: fulfilling the memex vision. In *ACM Multimedia* (2002), 235–238.

[6] Miller, G. A. Wordnet: a lexical database for english. *Communications of the ACM* (1995), 39–41.

[7] Rawassizadeh, R., Tomitsch, M., Wac, K., and Tjoa, A. M. Ubiqlog: a generic mobile phone-based life-log framework. *Pers Ubiquit Comput 17*, 4 (2013), 621–637.

[8] Wang, Y., Lin, J., Annavaram, M., Jacobson, Q. A., Hong, J., Krishnamachari, B., and Sadeh, N. A framework of energy efficient mobile sensing for automatic user state recognition. In *MobiSys '09*, ACM (2009), 179–192.