Personalizing System Behaviour in a Pervasive Social Networking System

E. Papadopoulou, N. K. Taylor and M.H.Williams
School of Math & Computer Sciences
Heriot-Watt University
Edinburgh EH14 4AS, UK
{E.Papadopoulou, N.K.Taylor, M.H.Williams}@hw.ac.uk

S. Gallacher
Intel Collaborative Research Institute
University College London
London, UK
S.Gallacher@ucl.ac.uk

***Abstract*—The aim of pervasive system technology is to support the user in managing and interacting with devices in the environment around him/her. By contrast social networking systems are aimed at supporting communication and interactions between people. The major objective of the SOCIETIES project is to unite these two technologies in a seamless manner to create a platform that enables the user to interact with both devices and people in an integrated way. Such a platform, referred to as a Pervasive Social Networking System, needs to integrate the pervasive computing technologies from the user's physical spaces with social computing technologies in the user's digital spaces. An important aspect of the SOCIETIES platform is its ability to adapt its behaviour to meet the individual needs and preferences of the user. This paper focuses on some key aspects of the design of the personalization subsystem of the platform: the extraction of personalization information from digital and physical spaces, the use of this to adapt these spaces autonomously and the extension to include community information to provide additional benefit to the user. Development of the basic platform together with a set of third party services running on it is complete and the system is currently being evaluated in a real user trial by a group of University students.**

Keywords—personalization; pervasive systems; social networking; user preferences; user trial.

I. INTRODUCTION

The ideas of mobile and pervasive computing have attracted much attention over the past decade. As the number of devices continues to escalate and the environment surrounding the user becomes ever more complex, the need for such systems grows increasingly important. A major challenge lies in creating an intelligent environment that provides support to the user in interacting with and managing these devices unobtrusively, without the user being aware of or needing to cope with underlying communications and computing technologies. However, the problem of capturing and incorporating the necessary intelligence to enable such systems to support users in different circumstances is non-trivial. Not only does such intelligence depend on the context of the user but it also needs to take account of the particular needs and preferences of the individual user which may be different for different users in different circumstances. Driven by this important challenge, research in this area has explored a range of different approaches based on different objectives, and a number of prototypes have been created to test these. Examples include Adaptive House [1], MavHome [2], GAIA [3], Synapse [4], Mobilife [5], Daidalos [6], Ubisec [7], etc.

By contrast social networking is a technology that has come into its own in a very short space of time. Systems such as Facebook, LinkedIn, MySpace, Bebo, YouTube, Flickr, etc., have opened up a whole new world of opportunities for communication via the Internet and, in so doing, have had a significant effect on social connectivity between users. From the extent to which such services are used on a daily basis by a significant proportion of users and the amount of time that users spend on them, it is clear that they have transformed the way in which computers are used by many users today.

However, these two technologies, pervasive computing and social networking, are complementary in many respects. It is our belief that if they can be combined and integrated seamlessly into a single system, there are significant benefits to be gained. Such a system, referred to as a Pervasive Social Networking system (PSN), would have the advantages of both pervasive systems and social networks together with benefits arising from the combination of the two.

The SOCIETIES project [8] has been building on recent technical developments in pervasive computing and social networking to create a prototype PSN platform. This was designed as a general platform with a view to testing it in several very different situations. In particular, the needs of three different trial groups were taken into account. These were:

(1) Students. A group of University students was recruited to provide input on their needs for such a system.

(2) Disaster Management. A group of international disaster management experts provided input on the way in which such systems could be used for disaster management.

(3) Enterprise. Input was obtained from our industrial partners, led by Intel, on the needs for such a system in an Enterprise setting.

The basic platform is now complete and a set of third party services developed to run on it. The system is currently being evaluated in a major trial which began on 23rd October when the system was set up to run for a group of twenty students at Heriot-Watt University and handed over to them to use.

One important characteristic of a PSN system is the ability to adapt its functionality and behaviour in accordance with the circumstances of the user, responding to changes in the environment as well as meeting the needs and preferences of individual users. In other words, a PSN needs to be context-aware and personalizable. However, in a PSN personalization may extend beyond the preferences of the individual to include preferences associated with communities. This paper is concerned with how a PSN deals with personalization in this broader context. In particular, the focus is on the SOCIETIES platform and how it deals with this problem.

The next section provides a brief background to the research while Section 3 describes the personalization subsystem used in the SOCIETIES platform. Section 4 describes the approach used to handle user personalization and how this has been extended to handle community personalization. Section 5 concludes.

II. Background

As mentioned in the previous section, much research has been devoted to the development of pervasive systems and a large number of prototypes have been created, based on different assumptions, using different approaches and aimed at different objectives. Endres, Butz and MacWilliams [9] give an overview of 29 software infrastructures and frameworks arising from different research projects in this area.

Much of the research has focused on the creation of systems based in a fixed location – or fixed smart space - of which the most important example is the Smart Home. Such systems are responsible for providing automatic control of devices supplying lighting, temperature control, security, communication, etc., and for balancing the needs and comfort of the occupants against factors such as energy conservation and safety. Examples of systems of this type include the Adaptive House [1], MavHome [2], GAIA [3], Synapse [4], Ubisec [7], the Intelligent Home [10], etc. Other important examples of fixed smart spaces include the smart office and smart buildings for large organizations or public buildings An example of the latter is Project Oxygen [11] which uses collections of embedded devices to “create intelligent spaces inside offices, buildings, homes and vehicles”.

Besides the work carried out on developing systems for fixed locations, there have also been systems developed to support mobile users. A number of prototypes have been developed to explore different pervasive system architectures and approaches for the mobile user, and to demonstrate and evaluate these. Examples include Daidalos [6], Mobilife [5] which focused on issues of privacy and trust as well as on maintaining a “shared cognition” amongst groups of users, Spice [12], etc.

The Persist project [13] provided a novel approach which spanned the gap between these two types of system. This was based on the notion of a Personal Smart Space (PSS). This hybrid approach covers both the fixed smart space (controlling sensor equipped buildings) as well as a mobile smart space that interacts with other surrounding fixed and mobile smart spaces.

In all of these systems one of the major goals has been to adapt the environment surrounding a user to provide an enhanced user experience. This relies on two key elements:

(1) Context. Context-awareness has been a major factor for these systems. In the simplest example, if a user enters a darkened room, the system might switch on lights, draw curtains, adjust temperature control, etc.

(2) User Preferences. In addition to the context of the user, it is important to know the needs and preferences of individual users. For example, simply switching on lights is not sufficient to guarantee an enhanced user experience. Some users may prefer full lighting, others subdued lighting and, if the user is blind, he/she may not care at all.

In some early projects such as Microsoft’s EasyLiving project [14] the collection of context information to drive adaptation of the environment was a major focus. This idea was taken a step further by Choi and Shin [15] who included biometric information to assist in decisions on environmental control. However, the need to capture information about the preferences of the individual user and use these to personalize the decisions taken by the system was soon realized. By personalization we mean tailoring the behaviour of the system to the needs of the individual in the particular context, thereby giving each user an enhanced and personal experience.

The user information of particular relevance to personalization comprises two forms: user preferences and user intent. User preferences consist of the knowledge (whether in the form of rules, neural nets, or other forms) for user specific adaptation, user intent consists of sequential patterns of user behaviour. Both forms are dependent on context. In both cases information could be entered manually, but in general needs to be automatically learnt by the system.

Learning of such user information is based on monitoring the actions performed by the user and applying machine learning techniques to the data. This approach has been used to create and update the set of preferences in a number of systems, including GAIA, Daidalos, Persist, etc. It has also been used to learn user task (user intent) information in the case of MavHome. Bayesian networks were used in the Synapse project for a similar purpose.

The choice of machine learning technique is an important one. On the one hand one has network approaches (neural nets, Bayesian networks) which cannot be checked or altered by the user, on the other hand symbolic techniques (generally some form of rules), the results of which are open to inspection by the user. One also needs to distinguish between batch learning algorithms that are applied infrequently to large sets of data and incremental algorithms which rapidly adapt to new information.

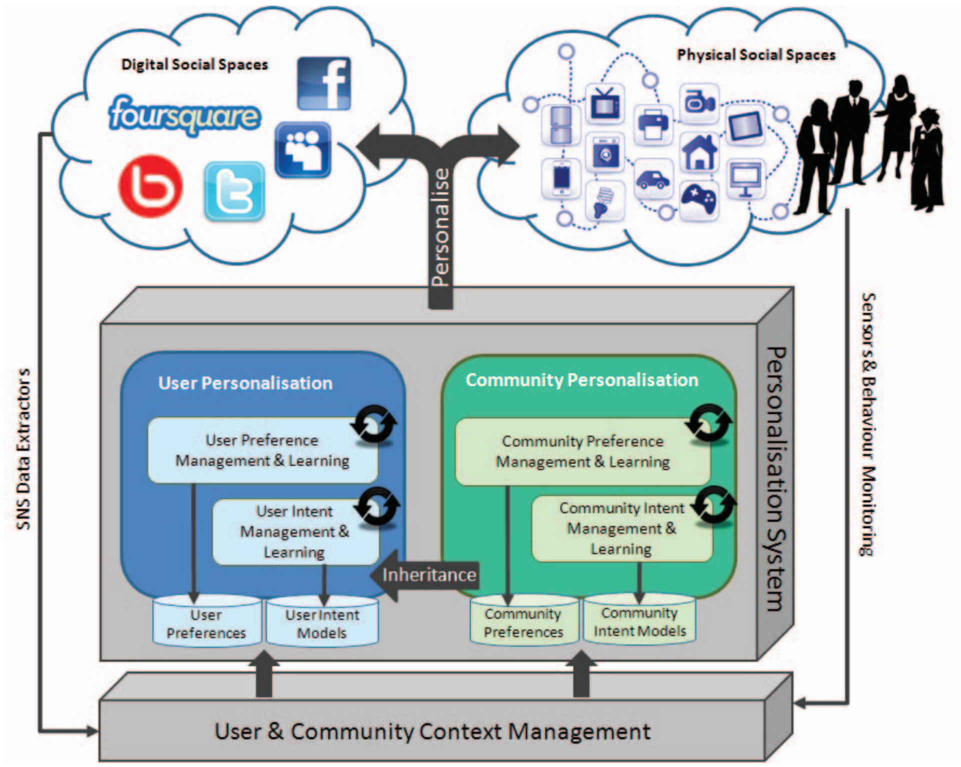

Fig. 1. SOCIETIES Personalization System

Nevertheless the task of learning user preferences (or intent) is highly challenging. Since most user preferences are context-dependent, building up a complete picture of a particular preference in all possible contexts in some cases will be impossible. Moreover, some (many?) of the preferences which a user may have, will change with time. Thus before the system can build a complete picture of the preference, it may have changed. In general it would seem to be impossible to have a complete and accurate set of preferences at any point in time. Accepting this, the challenge becomes one of obtaining a subset of user preferences that is sufficiently accurate to provide a level of adaptive behaviour that is acceptable to the user.

III. The Personalization System in SOCIETIES

In seeking to combine pervasive systems with social networking, this has a significant effect on the user preferences and user intent used to personalise the system. In a PSN we have the notion of a community, which is a collection of users who agree to share information and/or services between them. Any user can create a community and, depending on the membership conditions, other users may be invited to join the community. When a user wishes to join, he/she is asked to agree to the conditions attached to membership – these are usually concerned with agreeing to the sharing of information and/or services.

To illustrate the effect of this on personalization, consider the following scenario:

John is on his first day at University, having just registered for a course in Computer Science. He has joined the "Computer Science students" community. At lunch time his SOCIETIES device suggests a lunch venue on campus based on his physical location, his food preferences and the preferred venue of the Computer Science students community.

This scenario shows how the use of context (in this case location) combined with the individual's user preferences (food preferences) and the preferences of a community (the preferred lunch venue of the Computer Science students community) are used to suggest an appropriate venue for lunch.

To handle this, the architecture of the Personalization System within the SOCIETIES platform must deal with both individual user preferences and community preferences. A high level view of this is given in Fig. 1. This shows how the Personalization System is closely linked to the Context Management System for both User and Community contexts. The decisions produced by the system are used to personalize the user's interactions with devices in the physical social spaces as well as the social networking systems in the digital social spaces.

It also illustrates how the Personalization system can be regarded as two separate subsystems: User Personalization (UP) and Community Personalization (CP). The UP subsystem acts on behalf of the individual user, managing the user preferences and user intent models associated with the individual, the CP subsystem acts on behalf of communities of users.

IV. Extending User Personalization to Community Personalization

The User and Community Context Management System (UCCMS) is responsible for storing the context information associated with a user as well as the context information relating to communities. Sensors in the environment capture data such as location, proximity, temperature, etc., and pass this to the UCCMS where it may be refined (by inference rules or fusion processes) before storing it as context data.

Besides this sensor data, the actions of the user in relation to services is monitored in order to capture service usage behaviour. Every interaction between the user and a service is monitored, and the appropriate information concerning the interaction is stored together with a snapshot of the user's current context. This data set is referred to as the User Behaviour History (UBH).

Two separate approaches are used for handling individual user preferences within SOCIETIES – one based on a rule-based approach, the other on neural networks. In the former the system accumulates data on user behaviour and processes them in batches to generate a set of user preference rules, in the latter preference information is stored in neural networks and updates are handled on a continuous basis. Both are used to predict actions that the system should take on the user's behalf. If they disagree, a conflict resolution mechanism is invoked to choose between them.

Within the UP subsystem the batch processing is performed at an appropriate point in time on the back-end server. When this occurs, the system processes the user's UBH and extracts user preferences and user intent models. These are then used to update the existing set of preferences and models. The system continues to monitor user context and actions, and when it finds a match with either preferences or intent models, it will trigger the appropriate personalization actions.

The CP subsystem uses only the rule-based approach to represent user preferences. In this case the identification of new community preferences and intent models takes place at an appropriate point during the night. Each community is treated separately and analyzed for changes to the preferences and intent models. The new information arising from this process can be used in two ways. Existing members of the community can opt to inherit all or parts of the resulting community preference or intent data sets to enhance their own data sets. New members joining a community can inherit all or parts of the community personalization data, thereby giving them initial preference and intent sets. This is similar to the idea of using stereotypes to reduce the problem of building up such information from scratch.

In the case of any conflict between the individual's personalization and that produced by the community personalization data, the former will always take precedence over the latter. In particular, any conflicts between an individual's profile and that of a community that are identified during the inheritance process are handled with a bias towards the individual's profile.

However, one major issue with regard to the CP subsystem concerns how the community personalization data is obtained. Two possible approaches were identified, namely:

(1) For each community the system fetches the UBH of each individual user in the community, subject to the privacy constraints placed on this by the user. These individual UBHs can then be merged to create a Community Behaviour History (CBH). This data set can then be processed in the same way as the UP system processes the UBH, using machine learning algorithms to identify patterns in the data, and hence obtain community preferences and community intent models.

(2) For each community the system fetches the user preference set and intent models for each member in the community. It then performs a matching operation to identify any patterns that are common to all (or a majority) of the members. These then become the community preferences and intent models for that community.

After initially considering the first approach, we have now opted for the second in view of the potential size of the CBH for large communities.

V. Conclusion

This paper is concerned with the challenge of combining pervasive computing with social networking to create a new type of system, referred to as a Pervasive Social Networking (PSN) System. The focus of this paper is on the problems around handling personalization in such a system and the approaches used in the SOCIETIES project to overcome these.

The system as a whole has been implemented and is currently being evaluated in a real user trial by a group of twenty Computer Science students at Heriot-Watt University. They will be using the system for a period of six weeks. This will provide valuable feedback on the usefulness of community preferences and the problems that arise from them. We will be reporting on the results obtained in the first part of 2014.

Acknowledgment

This work is supported by the European Union under the FP7 programme (Societies project) which the authors gratefully acknowledge. The authors also wish to thank all colleagues in the Societies project. However, it should be noted that this paper expresses the authors' personal views, which are not necessarily those of the Societies consortium. Apart from funding the Societies project, the European Commission has no responsibility for the content of this paper.

References

[1] M.C. Mozer, "Lessons from an Adaptive House," in Smart Environments: Technologies, protocols and applications. D. Cook and R. Das (eds.), 2004, pp. 273-294.

[2] M.G. Youngblood, L.B. Holder, and D.J. Cook, "Managing Adaptive Versatile Environments," in Proc. 3rd IEEE Int. Conf on Pervasive Computing and Communications (PerCom '05), 2005, pp. 351-360.

[3] M. Roman, C.K. Hess, R. Cerqueira, A. Ranganathan, R.H. Campbell, and K. Nahrstedt, "Gaia: A middleware infrastructure to enable smart spaces," IEEE Pervasive Computing, vol. 1, pp. 74-83, 2002.

[4] H. Si, Y. Kawahara, H. Morikawa, and T. Aoyama, "A stochastic approach for creating context aware services based on context histories in smart Home," in Proc. 1st International Workshop on Exploiting Context Histories in Smart Environments, 3rd Int Conf on Pervasive Computing (Pervasive 2005), 2005, pp. 37-41.

[5] M. Strutterer, O. Coutand, O. Droegehorn, and K. David, "Managing and Delivering Context-Dependent User Preferences in Ubiquitous Computing Environments," in Proc. 2007 International Symposium on Applications and the Internet Workshops (SAINTW '07), 2007, pp. 4 – 4.

[6] M.H. Williams, N.K. Taylor, I. Roussaki, P. Robertson, B. Farshchian, and K. Doolin, "Developing a Pervasive System for a Mobile Environment," in eChallenges 2006 – Exploiting the Knowledge Economy, IOS Press, 2006, pp. 1695 – 1702.

[7] J. Groppe, and W. Mueller, "Profile Management Technology for Smart Customizations in Private Home Applications," in Proc. 16th Int. Workshop on Database and Expert Systems Applications (DEXA'05), 2005, pp. 226-230.

[8] S. Gallacher, E. Papadopoulou, N.K. Taylor, F.R. Blackmun, and M.H. Williams, "Intelligent Systems that Combine Pervasive Computing and Social Networking," in Proc. Ninth International Conference on Ubiquitous Intelligence and Computing (IEEE UIC 2012), Fukuoka, Japan, IEEE Computer Society, 2012, pp. 151-158.

[9] C. Endres, A. Butz, and A. MacWilliams, "A survey of software infrastructures and frameworks for ubiquitous computing," Mobile Information Systems Journal, vol. 1(1), pp. 41-80, 2005.

[10] V. Lesser, M. Atighetchi, B. Benyo, B. Horling, A. Raja, R. Vincent, T. Wagner, P. Xuan, and S. Zhang, "XQ.: The Intelligent Home Testbed," in Proc. Anatomy Control Software Workshop (Autonomous Agent Workshop), 1999, pp. 291 – 298.

[11] L. Rudolph, "Project Oxygen: pervasive, human-centric computing—an initial experience," in Proc. 13th Int. Conf. on Advanced Information Systems Engineering, Springer, 2001, pp. 1–12.

[12] C. Cordier, F. Carrez, H. Van Kranenburg, C. Licciardi, J. Van der Meer, A. Spedalieri, J.P. Le Rouzic, and J. Zoric, "Addressing the Challenges of Beyond 3G Service Delivery: the SPICE Service Platform," in Proc. Workshop on Applications and Services in Wireless Networks (ASWN '06), 2006.

[13] M. Crotty, N. Taylor, H. Williams, K. Frank, I. Roussaki, and M. Roddy, "A Pervasive Environment Based on Personal Self-Improving Smart Spaces," in Proc. Workshop on Constructing Ambient Intelligence, AmI 2008, Springer Verlag CCIS, vol. 32, 2010, pp. 58-62.

[14] B. Brummit and S. Shafer, "Better Living Through Geometry," J. Ubiquitous Computing, vol. 5(1), pp. 42-45, 2001.

[15] J. Choi and D. Shin, "Intelligent pervasive middleware based on biometrics," in Proc. Third International Conference on Ubiquitous Intelligence and Computing (UIC 06), LNCS 4159, Springer-Verlag, 2006, pp. 157-165.