Lifelong User Modelling Goals, Issues and Challenges

Judy Kay and Bob Kummerfeld

CHAI, Computer Human Adapted Interaction group, School of Information Technologies, University of Sydney, Sydney, NSW 2006, Australia
{judy, bob}@it.usyd.edu.au

Abstract. The combination of low and falling costs of storage, combined with increasingly pervasive and ubiquitous computing are creating the possibility of very long term user modelling. This paper explores some of the fundamental questions about the nature of a lifelong user model. It begins by exploring the relationship between the new goals that lifelong user models can support and the concomitant issues and challenges. We then explore the nature of what belongs in the user model, and what does not. From these foundations of the goals that the lifelong user model must serve, we discuss drivers for the design of systems which implement lifelong user models. We conclude with discussion of broad issues that follow from our analysis, proposing a roadmap for research needed to achieve lifelong user models.

1 Introduction

We are approaching a time when everyone will be able to inexpensively store large amounts of personal digital information. This may be distributed over a range of devices, including personal desktops, increasingly powerful carried devices, as well as in the cloud on various servers which can be accessed from a range of devices. In parallel, emerging pervasive computing will provide many different devices within the environment. Some of these will be appliance computers that perform restricted sets of tasks, just as a kitchen appliance like a toaster has limited functionality. Some of these pervasive appliances will provide personalised services. Others will provide interfaces that enable people to access and use their digital information: for example, a tabletop interface in a coffee shop or meeting room can provide the means for people to discuss projects aided by relevant documents that they share and work with on at the table. By exploiting the potential of inexpensive, ubiquitous and pervasive computing, it will be possible to provide valuable, extensive personalised services.

Consider this example of the way that a lifelong user model might operate.

Ted is is a retired concert violinist who has a long term interest in technology. He is 80 years old. Ted knows that he tends to forget things and that he takes longer to learn new things. He has some on-going health problems and recently injured his leg, becoming rather inactive.

Ted visits his GP for the results of recent medical tests. The GP explains the implications of these. Ted is somewhat flustered and concerned he will forget what the GP said. So, using the interface on his smart phone, he downloads digital information sheets and other notes provided by his GP. The doctor also recommends a new coaching programme for increasing exercise levels and fitness in the elderly. This too is loaded to his phone.

This scenario calls for several forms of long term user modelling. First consider Ted's long term goal of building up his fitness and exercise levels. This is typical of many important goals that remain relevant for people for long periods in their lives. Ted will work on this, aided by his new coaching programme. This will make use of the data about his activity levels. It will also maintain a model of Ted's knowledge about fitness and exercise and this can drive personalisation. It may also draw on older models of Ted's sporting activity and interests to tailor its recommendations. Some parts of these aspects of his user model may have been used by previous coaching systems Ted used. Ted may view these user models to see progress in improving his activity levels and fitness. Similarly, a visualisation of his knowledge in this area can help him monitor his progress and see the aspects still to learn. He may share these models with other people, such as a friend who is an exercise partner.

Now consider the other aspects of this consultation, where Ted learns about results of recent medical tests and collects some digital documents, information sheets and other notes. As Ted is forgetful, his user model may represent this, driving his phone to alert him to review these materials later in the evening. It may also remind him to show them to his daughter when he meets her the next day. His user model can include models associated with each of the documents from the consultation. These models include the context associated with his receiving them: the time, and place, the people present (the doctor) and links to associated documents such as the other documents received at that time, e.g. the test results. Each document model keeps details of the context each time he viewed it. These models can support *refinding* of documents, based on any aspects of context that Ted remembers.

2 Goals which need lifelong user models

The current state-of-the-art in user modelling includes considerable achievements in exploration of representations and approaches to reasoning as well as diverse ways to build user models, both directly via suitable user interfaces and automatically via machine learning and exploration of systems issues such as user modelling servers However, user models are generally restricted to a single application, often a single session. Some counter-examples include the work on user model servers [1,17,2]. Even so, user modelling research has been dominated by systems that have been evaluated in studies where each user made use of the system over relatively short periods of time.

Are these approaches sufficient for lifelong user modelling? To explore this, we identify some of the important classes of goals for which the lifelong user model is needed and their associated issues and challenges. Table 1 summarises these. We now discuss each in turn.

Achieving long term goals. This first goal must support personalisation over the long term, periods of several years. personal information management and personalised teaching systems. It is also needed for the personalisation envisaged for pervasive computing where systems integrated into the user's carried devices and environment provide useful alerts as well as "calm" displays of useful information. For such systems to operate, they must be able to query the user model for aspects such as their long term goals, knowledge, preferences, interests and attributes. They may also need short term contextual models such as the user's current task, activity and location.

Table 1. Goals for the lifelong model and associated issues and challenges

Purpose of user model	**Issues and challenges**
Achieving long term goals	How to deal with old models? Effective modelling of forgetting. How to take account of focus? How to resolve model values in differing contexts?
Reuse by different applications	How to deal with agreed meanings for the ontology used in different parts of the model? Whether we need to do this? How to address different standards defining model values?
User self-monitoring, reflection, planning	How to create effective user interfaces? How to provide meaningful model interpretations and values? How to support effective personal data mining?
User control and sharing	What representation can support effective interfaces? What interfaces will really enable user control?
Personal Information Management and Lifelogging	Should digital artifacts be regarded as implicit user models?
Institutional monitoring, reflection, planning	How to support the pragmatics of sharing parts of the model? how to support heterogenous data mining across these?

Returning to our scenario, it refers to three of Ted's long term goals: acquiring expertise as a violinist; a long term interest in technology; and maintenance and improvement of his health. The first of these is representative of the important class of learning toward outstanding expertise. Such learning achievement demands long term *deliberate practice*, "a highly structured activity, the explicit goal of which is to improve performance. Specific tasks are invented to overcome weaknesses, and performance is carefully monitored to provide cues for ways to improve it further" [5]. Key characteristics of deliberate practice are its long term nature "over a period of at least 10 years" [6]:368, that it is "not inherently motivating" but rather the learner does it to achieve "improvements in performance" and it is “an effortful activity that can be sustained only for a limited time each day” which is why it takes many years. Also notable is that it appears to require excellent teaching or coaching [6]. This class of learning might be supported by a long term user model which represents long term progress and activity and which supports teaching by a system or person as well as assisting the learner in monitoring their own progress and the effectiveness of their teachers.

Suppose that a single system helped Ted achieve expertise in one of these long term skills. For example, a system which teaches a foreign language, such as [7] may be used over long periods of time, with the learner returning to it as they have time or the need to revise or build their language skills. What are the new issues for *long term* user modelling in such situations? As shown in the table, there is a need for new work for the interpretation of old models. For example, if the user returns to the teaching system after 3 months, or a year, or five years, how should the old learner model be interpreted? How does this depend upon the level that the learner had achieved? The table shows that there are two forms of forgetting to be addressed. One relates to the user, as just described. The other refers to the model itself. Should the system maintain very old modelling information, dating back many years? With low cost storage, perhaps it is easy to simply keep all the information in case it one day turns out to be useful. But bloated models with information that is never used may incur problems in terms of both human issues of privacy, control and understanding of the model as well as technical aspects of scalability. These might be addressed by relegating older information to a different level of the storage hierarchy.

An associated problem relates to issues of focus. It is likely that the most recent modelling information is the most useful. This relies on the user's recent knowledge, preferences and interests. However, there may be other parts of the model that deserve to be treated as important. For example, suppose that a student used a single example repeatedly over a semester long programming subject. This could usefully be kept in an episodic user model [18] and it may be extremely useful to refer to this, even years later. It seems likely that heavy use of an artifact might usefully make it *stickier*, remaining as a focus point in the long term user model. However, other approaches may also deserve exploration. For example, mimicking human episodic memory [4], the user model might represent the richness of semantic connections to a particular concept. This may prove a useful indicator of its importance and that part of the model may be treated differently for modelling forgetting as well as importance.

Reuse by different applications. As longer term user modelling serves long term goals, there is an increasingly compelling argument for reuse of the user model. This introduces additional challenges, shown in the second row of the table. Because multiple programs may need to be able to make use of the same parts of the model, each must be able to interpret the model appropriately. This involves ontological challenges as well as the difficulties of potentially different standards and ways to assess the value of parts of the model. Consider an example of learning a complex skill, such as programming. One aspect is knowledge of loops. The term *loops* has an agreed meaning and it is easy to map this to other synonyms, such as *repetition* and *iteration*. However, the definition of when a user *knows* loops is more complex. In a teaching system for a first programming course, it may be reasonable to judge a user as knowing loops if they can read and write simple code in one language [10]. For a higher level student, one would expect far more, indicated by the ability to read and write more complex loop code, in more languages, and more quickly. A teaching system for senior programming students could reuse the model from the earlier stages. But it would need to "understand" several aspects beyond the semantics of *loops*: it should take account of the time since the student learnt about them, the source of the evidence about their knowledge and then it would need to interpret these. For some aspects, there may be established competency standards. However, it may be simpler and more accurate to set a diagnostic task on an as-needed basis.

User self-monitoring, reflection, planning. This has been the focus of work on Open Learner Models (OLMs) [3], also called transparent or scrutable models. The SMILI model [3] identified several key reasons for making a user model open to learners: enabling correction or verification of the accuracy; supporting reflection, planning and monitoring learning; providing an aid for collaboration and a basis for competition with others; as a guide to navigation, as in ELM-ART [18] where the colour-coded learner model indicated recommendations for aspects to explore (green), to avoid (red) as well as marking those completed (black); meeting the learner's right to access their own personal information and, finally, for assessment. All but the last are equally relevant for broader user modelling.

One of the essential elements for these goals is a suitable user interface to the user model. This can build on research in *Learner Modeling for Reflection* (LeMoRe [1]) which has explored many approaches to presenting user models and supporting interaction with them. Notably, if such an interface is created for a personalised system, then its user model, in association with the interface, can be of long term value for reflection, planning and monitoring goals. This can side-step ontological problems required for different programs to make use of a long term user model.

However, there is still the challenge of enabling the user to meaningfully interpret their model. For the important case of learning, we have already discussed the difficulties for a program to determine when something is well enough known. For the case of the individual learner, competency assessments or mastery standards may be useful. The more usual approach is to provide the learner with comparisons of their own performance against a suitable peer group. For well callibrated tests, the peer group may be based on age or educational level. More commonly, it is simply the comparison against the students enrolled in the same subject. For long term models, this may be of decaying usefulness.

One other important possibility provided by long term user models is the creation of new insights based upon data mining on the long term user model.

[1] urlhttp://www.eee.bham.ac.uk/bull/lemore/

While there has been considerable work on data mining large amounts of *collective* data about people, such as [12], and there is an emerging research area in educational data mining (EDM) [15], there is potential value in mining *individual* user models, for example, sequence mining to identify interesting and important patterns of behaviour. Even in collaborative contexts, this can be valuable, such as sequential mining of interaction activity to identify patterns of effective leadership [14].

User control and sharing. As indicated in the table, the challenge is to create interfaces that can make it easy for the user to achieve this. A suitable underlying representation and a focus on simplifying the user model are likely to facilitate the creation of such interfaces.

Personal Information Management and Lifelogging. Lifelogging, where large amounts of information is collected about the user's activities [13] has the potential to support augmented cognition. Similarly, personal information management involves large stores. Both these aspects reflect the new possibilities for building user models from collections of personal digital information or using these *directly* as user models as in work on personalised search [16]. This would be a radical departure from classic *explicit* user models. An intermediate approach could involve creation of explicit user models associated with files, modelling their importance and relationships with other files. This goes beyond existing approaches of creating metadata for files and sophisticated indexing and searching mechanisms.

Institutional monitoring, reflection, planning. This is currently done in large scale commercial data mining, such as [12]. However, there is potential for these approaches to be usefully applied in other contexts, such as the classroom where a teacher might gain understanding of the teaching effectiveness, using Educational Data Mining [11]. The issues for enabling this relate to the pragmatics of privacy of personal data and effective controls for an individual to be able to share just the parts of their user model that they wish. There is also the need to support data mining from heterogenous sources.

3 Discussion and conclusions

One key question is whether the nature of lifelong user models should continue the approach of much user modelling research, where user models are explicit sets of beliefs about the user [9]. This means that low level, raw data is interpreted and processed to contribute to the user model. This approach to user modelling underlies work on generic user modelling shells [8] user model servers [1,17,2]. A quite different possibility is to consider all of a user's digital artifacts and logs of their activity as part of an implicit user model. This underlies work such as the personalised search [16]. A third possibility is that the user model continues to be a carefully crafted representation of the user. To address the needs for user control, this should be supplemented with explanation and control subsystems designed to support the user in scrutinising their user model, amending it and controlling its release and use by different application. The model could also maintain links to relevant raw data, including the user's digital artifacts. A fourth possibility is to augment file systems with user models of files to support management of privacy as well as assisting in refinding based on high level criteria such as user goals.

This paper has explored the nature of lifelong user modelling, in terms of some of the key goals it can support and the associated issues and challenges. We conclude that very long-lived and long-used applications are likely to have an important role and each is likely to need its user model. There may be some cases where models need to be shared and for these agreed ontologies and interpretations will be needed. We particularly note the importance of user interfaces and open models to support direct use of the model by the user, to support their own reflection, monitoring and planning as well as for privacy management. And it is important to explore the appropriate linkages between the user's raw data and digital information on the one hand, and their user model, on the other.

The issues that are particular to the long term nature of lifelong user modelling define a new user modelling research agenda. This needs to explore the technical and social aspects of old models, the merits of explicitly building modelling systems to forget information that could clutter the model unduly and may pose other undesirable risks. We need to move beyond the exploration of ontologies for mapping terms, to the interpretation of modelling evidence in different contexts and at different times. We need to revisit our existing definitions of the nature of a user model: must it be separate from raw evidence and digital artifacts. And we need to explore the ways that we can enable people to make better use of their own long term models, with effective ways to analyse, mine, visualise and interact with them. Finally, we need to find ways to address the long present but unsolved issues of user control and privacy.

Acknowledgments

This research was supported by Australian Research Council Projects DP0774532, DP0877665.

References

1. M Assad, D J Carmichael, J Kay, and B Kummerfeld. PersonisAD: Distributed, active, scrutable model framework for context-aware services. In *Proceedings of Pervasive 07, 5th International Conference on Pervasive Computing*, volume 4480 of *Lecture Notes in Computer Science*, pages 55–72. Springer, 2007.
2. P. Brusilovsky, S. Sosnovsky, and O. Shcherbinina. User modeling in a distributed e-learning architecture. *User Modeling 2005*, 3538:387, 2005.

3. S. Bull and J. Kay. Student Models that Invite the Learner In: The SMILI:() Open Learner Modelling Framework. *International Journal of Artificial Intelligence in Education*, 17(2):89–120, 2007.
4. F.I.M. Craik and E. Tulving. Depth of processing and the retention of words in episodic memory. *Cognitive Psychology: Key Readings*, page 296, 2004.
5. K.A. Ericsson, R.T. Krampe, and C. Tesch-Römer. The role of deliberate practice in the acquisition of expert performance. *Psychological Review*, 100:363–363, 1993.
6. K.A. Ericsson, M.J. Prietula, and E.T. Cokely. The making of an expert. *Harvard Business Review*, 85(7/8):114, 2007.
7. W.L. Johnson, H. Vilhjalmsson, and S. Marsella. Serious games for language learning: How much game, how much AI? *Artificial Intelligence in Education: Supporting Learning through Intelligent and Socially Informed Technology*, pages 306–313, 2005.
8. A. Kobsa. Generic User Modeling Systems. *User Modeling and User-Adapted Interaction*, 11(1):49–63, 2001.
9. A. Kobsa and W. Wahlster. *User models in dialog systems*. Springer-Verlag New York, Inc. New York, NY, USA, 1990.
10. R Lister, E S Adams, S Fitzgerald, W Fone, J Hamer, M Lindholm, R McCartney, J E Moström, K Sanders, O Seppälä, B Simon, and L Thomas. A multinational study of reading and tracing skills in novice programmers. *SIGCSE Bull.*, 36(4):119–150, 2004.
11. A. Merceron and K. Yacef. A Web-based Tutoring Tool with Mining Facilities to Improve Learning and Teaching. In *Proceedings of 11th International Conference on Artificial Intelligence in Education*, pages 201–208, 2003.
12. B. Mobasher, R. Cooley, and J. Srivastava. Automatic personalization based on Web usage mining. *Communications of the ACM*, 43(8):142–151, 2000.
13. P. Olivier, J. Wherton, and A. Monk. State-of-Science Review: SR-E25 Technology Solutions to Prevent Waste of Mental Capital.
14. Dilhan Perera, Judy Kay, Kalina Yacef, Irena Koprinska, and Osmar Zaiane. *Clustering and Sequential Pattern Mining of Online Collaborative Learning Data*, page to appear. IEEE Transactions on Knowledge and Data Engineering. IEEE, 2008.
15. C. Romero and S. Ventura. Educational data mining: A survey from 1995 to 2005. *Expert Systems with Applications*, 33(1):135–146, 2007.
16. J. Teevan, S.T. Dumais, and E. Horvitz. Personalizing search via automated analysis of interests and activities. In *Proceedings of the 28th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 449–456. ACM New York, NY, USA, 2005.
17. J. Vassileva, G. McCalla, and J. Greer. Multi-agent multi-user modeling in I-Help. *User Modeling and User-Adapted Interaction*, 13(1):179–210, 2003.
18. G. Weber and P. Brusilovsky. ELM-ART: An adaptive versatile system for Web-based instruction. *International Journal of Artificial Intelligence in Education*, 12(4):351–384, 2001.