Scrutability, User Control and Privacy for Distributed Personalization

Judy Kay
School of Information Technologies,
University of Sydney, Australia 2006
judy@it.usyd.edu.au

Bob Kummerfeld
School of Information Technologies,
University of Sydney, Australia 2006
bob@it.usyd.edu.au

ABSTRACT

We have designed and built Personis, a foundation for personalized services that ensure the user can maintain control at all levels: what goes into their user model, what is available to different services and how the model is managed and maintained. We outline the approach and architecture.

Author Keywords
Guides, instructions, author's kit, conference publications.

ACM Classification Keywords
H5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

INTRODUCTION

Personalization is becoming widespread in web applications. It is also a goal of research into pervasive computing and areas such as teaching systems that aim to provide a personal teacher. There is considerable appeal in being able to exploit the large amount of digital information that could potentially support such personalization. We believe that a serious barrier to this is that, at present, user models tend to be hidden and out of the user's access and control. Our Personis [1] user modeling framework aims to maintain and manage user models in ways that overcome this, by a combination of the design of the underlying representation, the mechanisms for building and using user models and creation of a collection of user interfaces to support user scrutiny and control of their user models and the ways they are used.

ARCHITECTURAL OVERVIEW

Our broad approach is illustrated in Figure 1. Arrows at the top represent *evidence flowing from sensors*. In conventional systems, these include traces of web page visits, answers in education software and other logs of software activity. In pervasive computing, sensors may track the user's location and activity. Our architecture places a layer between the sensors and the user model: this layer controls which sensors are allowed to contribute to the user model. From a privacy perspective, this controls what is remembered about the user.

This detailed view of the user model shows that the actual model may be distributed. Parts of it may be stored in different places: at the user's home, workplace or other places they spend time. Parts may also be held on portable devices, such as a mobile phone or PDA. Each of these partial models has the controls over input of evidence and output of views with resolved values as shown in the main diagram. In addition, we show the evidence is regularly updated to a master user model. In fact, some queries, at one partial model will actually be directed to the master model, with two layers of control applied, once at the point of the query and again at the master model.

Arrows flowing from the bottom represent user modeling information for personalization services. The layer between these and the user model is critical for privacy and control of personal information and the way that it is used. Our Personis user model representation has two main mechanisms for this.

The first, *resolvers*, restrict the classes of evidence made available. For example, I can restrict strangers to evidence I have explicitly placed in the model, so excluding any evidence from any unobtrusive observations. In the common ubiquitous computing case where location is modelled, the resolvers available to different users may return results of different granularity. For a person who is modelled as a close friend, the answer may be a fine-grained precise location. For others, blurring is achieved by use of a resolver which is restricted to a small set of values, such as at-work or not-available.

The other mechanism, the *view*, controls which components are released: for example, whether my physics teacher may see the model for my maths knowledge.

A critical element of the architecture is the scrutability interfaces shown at the right. These must provide an overview of the user model, enabling the user to then scrutinize any part of the user model and the associated control elements described above. They should enable the user to drill right down to the full set of details, the evidence, the operation of the resolvers, the access details associated with views. In general, this is a challenging user interface problem. We have been exploring interfaces that give an effective overview of up to 700 concepts from the same context at once [2].

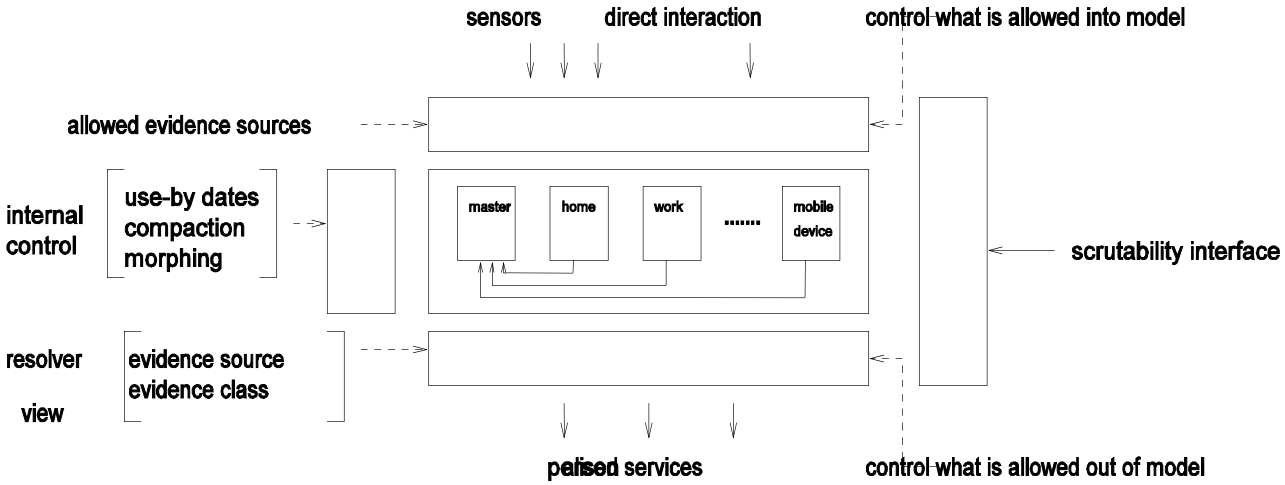

Figure 1. Architecture for privacy enhanced personalization

Finally, within the model, Personis supports privacy management: *use-by-dates* ensure that older evidence is removed; *compaction* replaces a set of evidence from a single source with one summarising piece; and *morphing* replaces an arbitrary collection. These processes reduce the size of the model. This, like the whole Personis design, aims to make models simpler and hence easier to understand and so to control the contents and use of the personal data that is the user model. The trade-off is that it may no longer be possible to make historic queries about the user model.

The underlying design of Personis has been strongly motivated by concerns for user privacy and control. In spite of this, effective control can only be achieved if we can create suitable interface. This is particularly challenging in ubiquitous computing environments where we have been exploring a P3P-like mechanism [3].

Privacy is an inherently personal issue, both in terms of personal preference about capture of information about a person and its subsequent use. Intuition, as well as studies of privacy preferences, clearly point to different people having different preferences and levels of concern about privacy.

On the management of information capture about a person, perhaps flowing into a user model, this means that we need to support flexible control mechanisms. The Personis approach is to characterise this in terms of the sources of the evidence. The user interfaces that can provide effective user control of this process need to ensure that users can, at the moment that they choose, work out how to manage this process.

The experiences with P3P suggest that this is challenging. Notably, this involves an adjunct user interface, making the usability demands very tough: users will make use of it infrequently. Part of the long term vision of Personis is that such information should be strongly linked to the user, with the personal data kept on the user's storage.

This somewhat simplifies management of inflowing user modelling evidence since it separates collection from the release. At the same time, it is subject to potential denial of service attacks by evidence sources that maliciously, or accidentally, produce huge amounts of evidence.

The management of information out of the user model is far more challenging. The Personis approach enables control at four main levels: the evidence source identity; the evidence source type; the component or component collection level including control via contexts and views; and the resolver level giving a particular interpretation of the available evidence. This is flexible, covering the main possibilities within the architecture. It will be challenging to create the interfaces enabling a user to express such preferences, and their combinations.

Overall, Personis represents one view of an underlying architecture that can give scrutability and associated control on the user models at the core of personalisation. Effective support for scrutability and control is currently particularly difficult because user's mental models do not include these notions: indeed, people may be attuned to expecting software to behave unpredictably and to having no means to scrutinize this to work out why, let along to control it. Progress will require both developing those mental models and fine interfaces to support scrutability and control of personalisation.

CITATIONS

1. Kay, J., Kummerfeld R.J. and Lauder, P., (2003) Managing private user models and shared personas, Cheverst, et al(eds), UM03 Workshop on User Modeling for Ubiquitous Computing

2. James Uther and Dr Judy Kay: Describing and Viewing Large User Models. *TR 522*. 1999. http://www.it.usyd.edu.au/research/tr/tr522.pdf

3. Ajay Brar, Judy Kay: Privacy and Security in Ubiquitous Personalized Applications. *TR 561*. December

2004. http://www.it.usyd.edu.au/research/tr/tr561.pdf