

# Personal Ontologies for feature selection in Intelligent Environment Visualisations

David J. Carmichael, Judy Kay, Bob Kummerfeld  
School of Information Technologies  
University of Sydney  
NSW 2006 Australia  
{dcarmich,judy,bob}@it.usyd.edu.au

## ABSTRACT

A central problem of Intelligent Environments is the *Invisibility Problem*. This arises when computing services are hidden in the environment in order to be unobtrusive and natural and yet the user is unable to make use of them because they are unaware of their existence. In this paper we describe a key aspect of the invisibility problem: the definition of *personal* models of places and services. We describe the use of Verified Concept Mapping to enable users to define personal ontologies of places and their associated people, sensors and services.

**Keywords:** Intelligent Environments, Invisible Information, User Modelling, World Modelling, Personal Ontologies

## 1. INTRODUCTION

One branch of Ubiquitous Computing aims to augment the physical environment with many computational devices and sensors to assist people in their activities. Such environments are called Intelligent Environments (IEs).

Unobtrusiveness and Invisibility are goals of good intelligent environments. Systems should be so natural and well hidden in the environment that unless a user knows what components are there, they are almost invisible. This invisibility also has a problem: if things are so well hidden, how does a new user find out what services are available to them and what sensors are detecting them. We call this *The Invisibility Problem*.

As an example consider Fred and Jenny, new postgraduate students commencing their degrees in a smart building. The building contains a number of invisible or non obvious systems such as bluetooth based location sensors, e-ink based notice boards, and wireless networking.

As Fred enters the building he passes a notice board displaying undergraduate notices which he ignores and keeps walking. He does not realise that, were he to identify himself as a postgraduate student, the notice board would present items relevant to postgraduates.

Jenny, on the other hand, enters the building carrying a bluetooth-enabled phone and is detected by the loca-

tion system. As she approaches the notice boards, they change to show information she might be interested in. She is somewhat disturbed by this as she doesn't know how the system knows she is there.

These two cases give examples of both parts of the invisibility problem. In the first example, Fred is unaware of the services available to him and in the second, Jenny is unaware of the sensors detecting her.

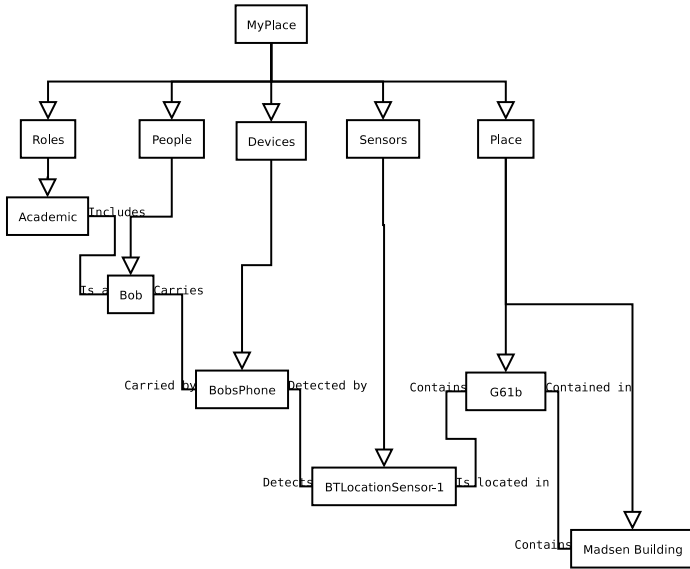
Related to the Invisibility problem, we must also consider the fact that such environments can be personalised. In the example above, the personalisations of the environment - tailoring the notice boards to the nearby users - is essentially a complicating factor. The fact that the notices were different for Fred and Jenny is yet another thing about the environment which must be explained to them.

However, personalisation can be useful if it is done more effectively. If the system models users more closely, it could reason that neither Fred or Jenny know about the functionality of the smart notice boards or the bluetooth location sensors and inform them.

In an IE there may be thousands of known places, devices, sensors and services. In order to present useful information to the user, we must develop ways to reason about what the user is interested in and already knows about. We want to provide personalised information about services and sensors in a personalised IE.

One key element to our work is the definition of personal models of places and their associated sensors and services. We tackle the personalisation of this with Verified Concept mapping to get the user to input what they know about and are interested in. Users are given a partially constructed concept map and can modify it to create a personal ontology of concepts within the Intelligent Environment.

The rest of the paper is organised as follows. Section 2 describes the MyPlace system. Section 3 describes Verified Concept Mapping and how we use it to get personal ontologies from the user. Section 4 describes related work. Section 5 contains a discussion of future work and conclusions.



**Figure 1: An extremely simplified world model showing the linkages between items**

## 2. MYPLACE

In this section we describe the MyPlace system [8] which provides personalised views of ubiquitous computing environments in order to alleviate issues such as the invisibility problem described above. The novel aspects of this system are the uniform representation of users, places, locations, devices and services, the double personalisation and the *accretion-resolution* used for the world model.

### 2.1 Uniform Representation

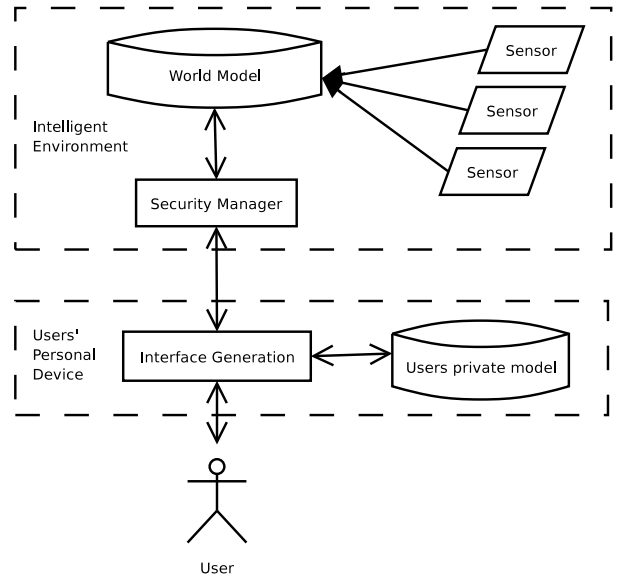
The uniform representation means that all entities in the system are organised into an inheritance hierarchy of concepts. General concepts form the higher level nodes and the entities modelled become more specific as one moves down the tree. In Figure 1 a particular room (G61b) is a child of the more general concept of Place. The nodes of this tree are also interconnected by their semantic relationships. In this example ‘BTLocationSensor-1’ is located in ‘G61b’.

Each concept can have simple valued attributes in addition to the semantic linkages to other entities. Concepts can have arbitrary attributes and linkages as there is no fixed schema.

### 2.2 Dual Personalisation

MyPlace [8] includes double personalisation: The content is personalised for the user in two stages. The architecture is shown in Figure 2. First the Security Manager generates a subset of its total world model based on what the user is allowed to use or know about. Second, the user’s personal device selects what parts of the restricted world model the user should see.

This model gives two major benefits. Firstly the IE can select what the user is allowed to see. For exam-



**Figure 2: The delivery architecture for MyPlace.**

ple, postgraduates are not allowed to use the expensive colour printer, and thus are not told about its existence. Secondly, the user is able to have preferences about what they care about without these preferences being released to the IE, thus protecting their privacy.

### 2.3 Accretion-Resolution

Another interesting aspect of MyPlace is the manner in which changing values are represented. As information is received from sensors in the system it is simply stored as evidence without interpretation - this is the *accretion* step. Then, when a value is requested, the stored evidence is examined and *resolved* to a value. For example in Figure 1, when BobsPhone is detected by BTLocationSensor-1, it is noted in the models for both entities as time stamped evidence. When we wish to know Bob’s location, all the evidence about what devices Bob is carrying and by what sensors they have been detected, is resolved to determine a value. The advantage of this approach is that it saves the system from having to put effort into interpreting sensor data as it enters the system.

## 3. VERIFIED CONCEPT MAPPING

One of the core challenges in defining personal ontologies of places involves user interface support for the creation of personal ontologies of places and their associated sensors and services. We tackle this with VCM [4], the verified concept mapper. This is based on a well established knowledge elicitation technique called concept mapping. Its origins are in education [11, 10, 9] where it was first used to enable learners to externalise their understanding of the relationships between the concepts in a learning area. It has been widely used in education and has also been used for knowledge elicitation[5].

A typical use of concept mapping, as we apply it, is to provide the user with a collection of concepts, some



**Figure 3: Screenshot from the MyPlace system for Fred, a male staff member in the school.**

of which have already been placed into a concept map. This defines triples in terms of pairs of concepts and the relationships between them. In addition to this obviously semantic element, concept maps also make use of the two-dimensional layout. The user places concepts that the user considers similar in a vertical line. Also, concepts further to the left in the map are high level, compared with those at their right. The user also strives to make the maps symmetrical and pleasing to the eye. Essentially, the physical location of concepts in the map implies some relationships even when they are not captured as semantic links. Moreover, the presentation is intended to be natural, making it easily understood and memorable.

These properties of concept mapping have contributed to its wide use and success in educational contexts. This is why we considered it an appealing foundation for our interface that would support users in amending and extended foundation ontologies of spaces with personal ontological differences. Its usefulness in education, even in the case of young children, suggests that it might be usable by a diverse range of users who want to understand and control the models that drive a personalised intelligent environment.

VCM extends basic concept mapping by verifying the maps created. Once the user indicates they have completed a map, the verification stage checks the map and provides two forms of response: suggestions for the user to consider and a summary of the inferences it has drawn from the map.

The suggestions are based upon two classes of checks. First, the author of the mapping task can define map features that they expect to be part of the map. In the context of the MyPlace work, authors would define the

foundation map and associated concepts and relations to be used in individual maps. If any of the expected features is missing in the user's personalised map, this is signalled to the user, either via a default message or a specialised one written by the author of the foundation map. For example, if all staff are expected to identify their own office and workplaces, omission of these would cause a message asking them to check this is what they intended. The second form of suggestion relates to unexpected features found in the map. This operates similarly, being coded by the authors of the foundation map, giving either default or specialised messages.

The second part of the verification phase is the list of amendments the user has made to the map. This lists new concepts as a triple defining the concept they are linked to and the relationship on that link. Since this would generally be a quite short list, it should be plausible for a user to check this. It also lists elements the user has deleted from the foundation map.

This list is a basis for subsequent explanations of differential personalised behaviour in MyPlace: if a person asks for explanations of the way MyPlace presents information, these will be expressed in terms of their own changes to the foundation ontology.

The essence of the verification is an interaction with the user to improve the likelihood that accidental errors are likely to be identified: it is very easy for a user to click on the wrong concept or link. Once the verification phase is complete, we treat the map as a reliable model of the user's ontology of the space. Moreover, the user is aware of the elements they have altered and have had their attention directed to unusual or unexpected elements that they have defined.

This section describes how the user interacts with VCM. The interaction in context Figure 3 shows a screenshot from MyPlace.

### 3.1 Modelling Spaces

Figure 4 shows the foundation ontology for one region using the VCM system. The panel on the left hand side shows all the concepts known to the system. Colour is used to show the category these concepts belong to: places (black and blue); functions (red) and personal activities or annotations (green). The semantics of the system-supplied concepts is known. The user can define their own terms but the system will not be able to associate semantics with user-defined terms. The right hand side gives a list of relationships the user may define between concepts. Again colour is used to give information about the meaning of the linkage. Geometric relationships such as Contains or Adjacent are shown in black, room functions (Is) are shown in red and terms for use in personal ontologies are shown in green.

The centre area of the screen is where the user lays out concepts and links them together in concept maps. They are provided with a foundation map of concepts which they can add to, modify or delete (although we do not show an example of deletion). It is critical to




Figure 4: The VCM system showing its foundation ontology.

provide the foundation map for two reasons. Firstly, this saves the user the tedium (and potential errors) in defining parts of the ontology that are generally common. This is defined by the architecture of the building and spaces as well as the location of sensors and services. This will generally change slowly. This is also important in providing consistent images of the core elements of MyPlace. Secondly, we want to ensure that the number of personalised elements in a MyPlace ontology will be small enough for a user to readily distinguish and check those elements.

Figure 4 shows the Locations known to MyPlace placed in default positions. This arrangement of places is done the same regardless of who the user is. Within the University of Sydney (denoted Usyd) there are two buildings which the system knows about, Madsen and Storie Dixon. Rooms within these buildings are labelled by their numbers: G70, G71, G92 ... Lab381a, Lab381b. The author defined the map only with places relevant to the current MyPlace implementation.

The next stage in the mapping process is the automatic labelling of rooms based on the type of user. Similar to Role Based Access Control, the user's access is determined by which groups they belong to. Figure 5 shows the concept map for Fred, a male staff member. We can see labels on all the shared rooms that a staff member

would be expected to know about. The Postgraduate workspace and undergraduate laboratory are shown as staff are expected to supervise research students and teach some undergraduate classes.

Figure 6 shows the foundation map generated for Jane, a new female postgraduate student. There are 2 differences from Fred's map, firstly the undergraduate laboratories are unlabelled as new postgraduates are not expected to do any teaching. Secondly instead of labelling the male toilets the female toilets are labelled.

The final automatically generated concept map is shown in Figure 7. This is for John, a visitor coming to give a guest seminar. In this case he is only given semantic labels on the Seminar Room and the Front Office.

To this point the concept maps are all foundational. They might be presented to a user as part of the explanation for what is displayed by the MyPlace interface of Figure 3. We now describe how the individual user can amend a concept map with their personal ontological information.

Figure 8 shows Fred's personal annotations. He considers himself in his office when in G61a and he does research and teaching in the places marked as MyLabs. When he is in the seminar room (G92) or the meeting room (G46) he expects to be in meetings. As a result of

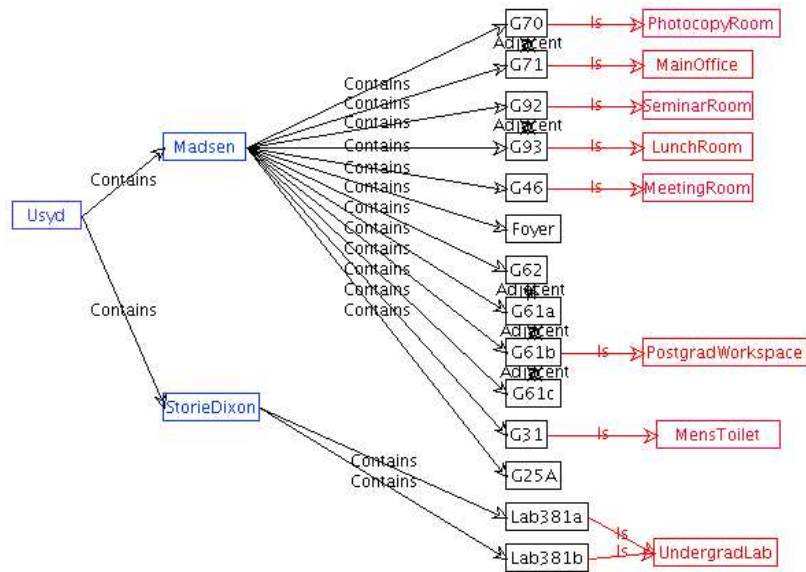


Figure 5: VCM showing annotations based on the type of user, in this case Fred is a male staff member.

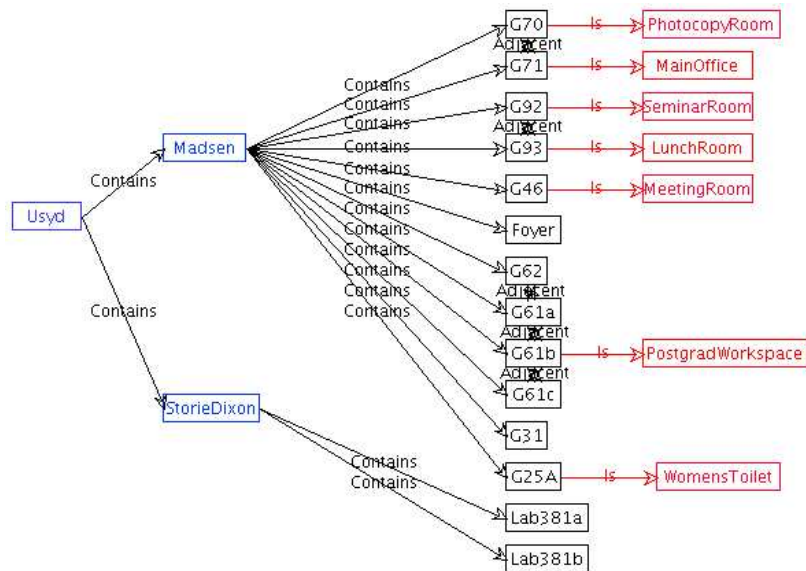
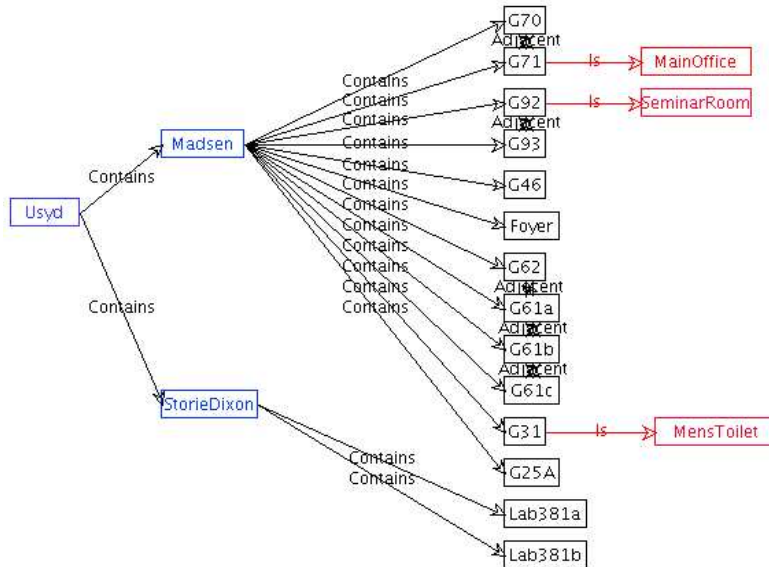


Figure 6: VCM showing annotations based on the type of user, in this case Jane is a new female postgraduate student.



**Figure 7: VCM showing annotations based on the type of user, in this case John, a visitor giving a seminar.**

being given this information, MyPlace can reason that in either of these locations Fred is uninterruptable as he is in meetings. (He can override this behaviour if he wishes, but that part of the system is outside the scope of this paper.) Fred has specified that he is to be considered at work when he is located in the more general location of the Madsen building rather than a specific room. This allows the system to aggregate location information and return the more general concept, "at work". MyPlace uses this in the reporting of Fred's location, eg user A is only informed if Fred is at work or not.

Jane annotated her diagram slightly differently in figure 9. Her office is G61a (the postgraduate workspace). She has linked the InMeetings concept to an additional room - G61a, her supervisor's office. She believes that when located in that office she will be discussing research with her supervisor and thus want the system to treat her as being in a meeting. She has also created a new term SupervisorsOffice, which was not on the provided list. The system will not be able to automatically reason based on the creation of this concept, but she can define rules using it to allow this later.

After the user has made their additions and changes to the concept map it is verified as was described in the introduction to Verified Concept Mapping. This allows the system to flag map features that were unexpected as well as potential omissions the user may have made. By consulting rules entered by the author of the foundation map, the system can detect anomalies such as users having meetings in the toilets or not choosing an office.

We have similar ontologies for people and devices which have been omitted for space reasons.

#### 4. RELATED WORK

There is a number of related projects which deal with some of the elements of MyPlace. We now briefly outline other work which deals with aspects of the invisibility problem, the modelling of users, devices, activities and places and the personalisation of information delivery in pervasive computing systems.

The Digiscope described in [12] presents a system where the user is given an augmented reality view of invisible information in the environment by looking through a movable large semi-transparent perspex window.

Chalmers et al.[2] have constructed a framework for *context mediation*, where they adapt the content of documents based on preferences about its semantic and syntactic properties.

At Dartmouth College researchers have worked on an event based system for processing context data by defining processing operators and connecting them in an acyclic graph[3]. In our work we are attempting to allow users to specify what they are interested in having displayed to them using a similar idea.

Other approaches to modelling the world have not taken a uniform approach. CMU's project Aura[7] models entities in the world as either devices, people, areas or networks. The NEXUS project[1] focuses more on location and has a much stricter schema.

Huhns and Stephens[6] believe that having people make personal ontologies could assist them with managing and finding their documents, both real and electronic.

#### 5. CONCLUSIONS AND FUTURE WORK

The MyPlace system currently works for users' personal devices within our department. We are currently work-

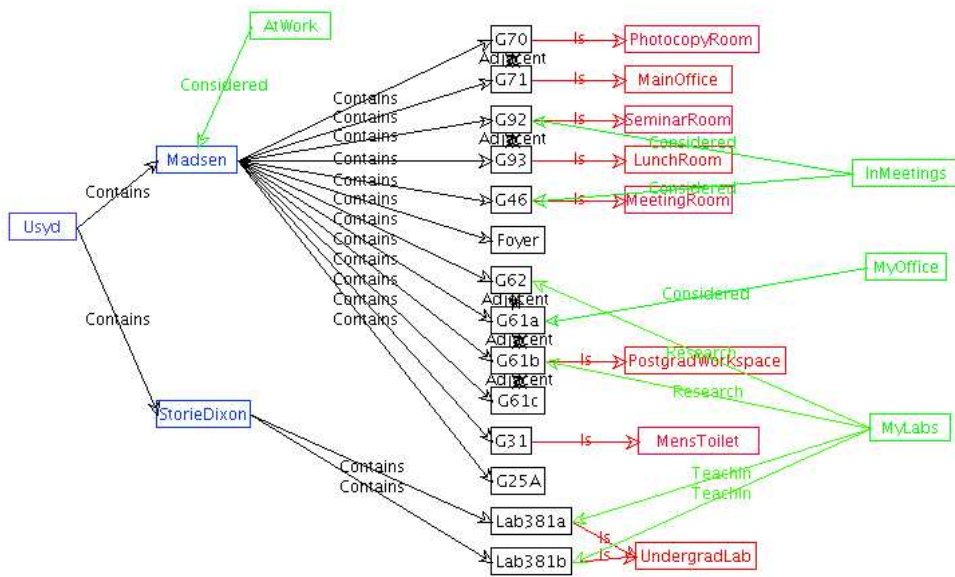


Figure 8: VCM showing the user's personal annotation. This map shows Fred's personal annotations, such where he considers himself to be at work, in meetings or in his office.

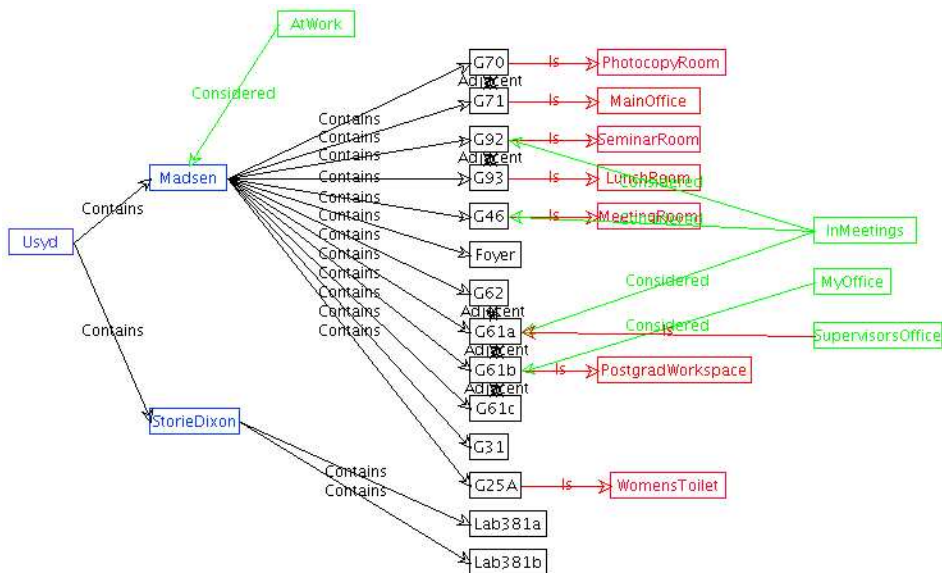


Figure 9: VCM showing the user's personal annotation. This map shows Jane's personal annotations, such where he considers herself to be at work, in meetings, her office or discussing research with her supervisor.

ing on ways for the user to further customise which elements of the world model they are interested in.

#### Acknowledgements

The authors would like to acknowledge the support of the Smart Internet Technology CRC.

#### REFERENCES

- [1] Martin Bauer, Christian Becker, and Kurt Rothermel. Location models from the perspective of context-aware applications and mobile ad hoc networks. *Personal and Ubiquitous Computing*, 6(5-6):322–328, December 2002.
- [2] Dan Chalmers, Naranker Dulay, and Morris Sloman. A framework for contextual mediation in mobile and ubiquitous computing applied to the context-aware adaptation of maps. *Personal and Ubiquitous Computing*, 8(1):1–18, 2004.
- [3] Guanling Chen and David Kotz. Context aggregation and dissemination in ubiquitous computing systems. In *in Proceedings of Forth IEEE Workshop on Mobile Computing Systems and Applications*, pages 105–114, June 2002.
- [4] Laurent Cimolino, Judy Kay, and Amanda Miller. (to appear) concept mapping for eliciting verified personal ontologies. *Special Issue on Concepts and Ontologies in WBES of the International Journal of Continuing Engineering Education and Lifelong Learning*, accepted April 2003, to appear, 2004.
- [5] Brian R. Gaines and Mildred L. G. Shaw. Knowledge acquisition and representation techniques in scholarly communication. In *Proceedings of the 13th annual international conference on Systems documentation*, pages 197–206. ACM Press, 1995.
- [6] Michael N. Huhns and Larry M. Stephens. Personal ontologies. *IEEE Internet Computing*, 3(5):85–87, September - October 1999.
- [7] G Judd and P Steenkiste. Providing contextual information to pervasive computing applications. In *Proceedings of IEEE International Conference on Pervasive Computing (PERCOM)*, pages 133–142, March 2003.
- [8] Judy Kay, Bob Kummerfeld, and David J Carmichael. Consistent modelling of users, devices and environments in a ubiquitous computing environment. Technical Report TR 547, School of Information Technologies, University of Sydney, <http://www.it.usyd.edu.au/research/tr/tr547.pdf>, June 2004.
- [9] J. D Novak. The theory underlying concept maps and how to construct them, <http://cmap.coginst.uwf.edu/info/> (visited july, 2004).
- [10] J. D Novak. *Learning, creating, and using knowledge: Concept maps as facilitative tools in schools and corporations*. Lawrence Erlbaum Associates, Mahweh, NJ, 1998.
- [11] J D Novak and D B Gowin. *Learning how to learn*. Cambridge University Press, 1984.
- [12] Alois Rerscha and Markus Keller. Digiscope: An invisible worlds window. In *Video in 5th International Conference on Ubiquitous Computing*, October 2003.