

# Personalized Cultural Heritage GeoNotes

Daniel Brem, Kurt Gubi, Judy Kay, Bob Kummerfeld, Tsvi Kuflik, Joseph Michaels, and Rainer Wasinger

University of Sydney, Sydney NSW 2006, Australia

**Abstract.** Geonotes related applications are booming. This has the potential to create an overload of geonotes - wherever we go, there may be tens, perhaps hundreds of geonotes. Even if all are informative and interesting, we simply cannot view them all; and many of them will not be of interest to a particular person. Cultural heritage is a domain where there is both the potential for valuable information delivery via geonotes. Personalization offers a way to reduce, perhaps avoid, this anticipated information overload. To make this work, we propose to enable the person creating a geonote to associate it with information to support personalization. This can include: a relevant part of their own profile; a model of the target users and the context for the geonote. Such personalized geonotes can be filtered on the set of such models leaving just those most relevant for them. This paper introduces the novel idea of personalized geonotes, using an extended scenario and outline for an architecture to support this.

**Keywords:** Personalization, Geonotes, Lifelong User Modelling

## 1 Introduction

Geonotes provide a mechanism for associating asynchronous messages with a location. This seems an appealing way to allow a person to augment the environment with information that could be helpful to others, or even for themselves at a later date. Early geonotes work [4] recognized the need for mechanisms to overcome the potential overload if all notes that anyone leaves at a place are equally available to all people who visit that spot. The extent of this problem is indicated in the rapid growth of the Foursquare geonote service, which had more than 500,000 users and 1.4 million venues within a year of its launch [7]. The focus of our work is to exploit information about the person who posted the geonote and about their target audience and other context to define the relevance of each geonote for each user.

Geonotes are particularly appealing as a means of providing rich personalized information about cultural heritage sites. It offers the promise of different notes with different information to meet the diverse needs and interests of different visitors. Consider the following four part scenario:

1. *Alice teaches third grade. She is preparing for her class visit the Australian Museum as part of their studies of spiders, to gather information for creating*

*a poster. The children are working in groups of 3-4, with each student allocated their one aspect; one child is studying the spiders that live in the school area; another on typical spider habitats and diet; a third on how they make people's lives better; and optionally a fourth on how dangerous they are and how to treat bites. Alice uses the Museum geonote authoring system to place geonotes from the class e-portfolio system, with each child's own poem about spiders linked to locations in the museum (to help them see the link between that poetry writing activity and this visit). Also, each child gets a set of the core information. But each also gets their information according to their role in the group work for the poster.*

*2. When the class visits the museum, each child uses a mobile phone to access their own notes in each part of the museum. They can augment the notes with their own, additional information, based on the materials in the museum. Later, each group comes to a tabletop display, releasing all their notes to the table so they can discuss them to assemble the information for the poster.*

*3. A month later, Bob, a teacher of a 4th grade class plans to bring his class to the museum to study spiders. They will do individual poems about spiders and how they help people. He finds Alice's collection of geonotes and decides to reuse some of them. He also creates some variants, extra geonotes and changes the personalization so that a selection of the notes are presented to every child in his class.*

*4. Later, David, a retired entomologist, visits to the museum. He discovers a geonote left by a child, during a class visit. He leaves some follow on geonotes that answer the questions posed in this note.*

This scenario relies on considerable infrastructure that goes beyond that available in widespread geonote systems. This paper presents a high level architecture to provide that capability. The next section summarizes related work. Then we describe the architecture and discuss the challenges of achieving it.

## 2 Related Work

This work builds on four main bodies of previous work. First, it is a form of geonote system and so we describe the previous work in that area and the state of the art in deployed systems. Second, we make use of an explicit user model for the poster and target audience of the notes and so we briefly outline some of the ways that personalization based on such a model has been used to deliver a personalized cultural heritage experience. The third aspect is to address the challenges of indoor location in a manner that matches the demands of delivering personalized geonotes. Then we briefly describe the foundations that school curricula can provide for an ontology that can be used to describe museum visitors and their interests over the long term.

There is considerable appeal in enabling people to create interesting material to associate with a place. For example, Pascoe [14] proposed the notion, inspired by the potential of location awareness applied to existing uses of physical stick-e notes. The Cyberguide series of prototypes [1] included forms of geonotes. These

could be simple text notes as in the early GeoNotes [4] system or audio as in a system like murmur <sup>1</sup>, which enables people to leave oral history stories or recollections at places marked with a sign, so that others can dial a phone number to hear the stories at that place. The particular way that a geonote system operates can influence the nature of its use. For example, the early GeoNotes system was designed to favor interpersonal interaction among friends and analysis of its use indicates that this use dominated

Being rich in information, cultural heritage sites have proved an appealing area for research in personalized delivery of information. There is always too little time and too much information about museum artifacts, and this means that a personal user model may help a system guide museum visitors more effectively. For example, PEACH [15] provided personalized and context aware multimedia presentations to museum visitors. It tracked visitors' interaction with the system and built an overlay user model (several different types were developed and experimented with), over a domain knowledge base. This model evolved during the visit and allowed tailoring of the information to the individual user's preferences. Moreover, keeping a record of the route in the visit, PEACH was then able to tailor the presentation, referring to things the visitor has already seen (or would see). Once the visit was over, visitors received a personalized visit summary that included suggestions for nearby attractions and future visits.

Another important example of personalized museum tours is the CHIP project [16] which combined activities before and during the museum visit. CHIP allowed users to view and mark artwork that they were interested in before the visit. This enabled the system to build a user model for them and to suggest more relevant pieces of artwork to view. Later on, when the user visited the museum, CHIP could provide a personal on-site museum tour based on the information accumulated in the user model, from the earlier online interaction.

The idea of personalized geonotes has already been partially explored in cultural heritage sites. Many museum and tourist guides allowed their users to send notes and leave virtual notes (geonotes) at specific physical places. The PIL project [12] allowed museum visitors to leave virtual *post-its* for other members of their group visiting the museum. The *post-its* had an intended audience (an individual or a group), a *time to live* (including forever) and context-aware content, commenting on a specific exhibit or specific presentation about an exhibit. CyberGuide [1] was a mobile, context aware tour guide that allowed a user to leave messages to exhibit owners and to send reports about their location to a central service, making them accessible to others.

Another early system was Guidebook [6], developed for the Exploratorium in San Francisco, which supported two communication functionalities: the *rememberer* enabled the visitors to capture a record of their experiences to consult during and after their visit; the *communicator* helped visitors communicate via electronic bulletin boards for individual exhibits, instant-messaging, and/or beaming information between handheld devices.

---

<sup>1</sup> [www.murmurtoronto.ca](http://www.murmurtoronto.ca)

As well as the *explicit* user input to tag the geonote’s location, we can also gather the locations a user has visited as *implicit* user input. This follows the work on modelling a *place* rather than a *location* as in BeaconPrint [9]. From the semantic information provided by geonotes we can automatically determine rich semantic descriptions of locations in terms of available *symbolic maps*, where these were designed to help the user see the parts of a complex space that are relevant to them.

The final foundation for this work involves the definition of the ontology that will be used to model the users. In our scenario, one promising approach is to exploit some of the rich existing, though implicit, ontologies that are represented in formal syllabus documents (such as [13]). Such documents define broad levels of development. In our scenario, the NSW syllabus would place Grade 3 students within a two-year grouping (Stage 2). This means we can model Grade 3 students using a *stereotype* that assume they have mastered Stage 1 aspects of the syllabus. More importantly, teachers and others involved in creating learning materials for students are familiar with the syllabus, its terms and the nature of what students learn in that school area. For students from elsewhere, for example tourists, assumptions based on this model are likely to be inaccurate but may still provide an inexpensive and useful set of default assumptions. A particular attraction of exploiting this approach is that it offers the potential for a relatively low cost approach to building comprehensive models, as a basis for more effective personalization.

### 3 Pigeon architecture

To support the forms of personalized geonotes of the above scenarios, we need to create an infrastructure that goes beyond that of the basic geonote system. Figure 1 gives a high level overview of the Pigeon architecture. It has three main elements. The upper right cloud is the Pigeon server. The teacher interface is shown at the lower right and the part of the system that resides on the student’s phone is at the left.

Before describing the details of the architecture, we need to define the information that is associated with each geonote:

1. the content of the geonote itself;
2. the location associated with the note;
3. additional context;
4. a partial user model for the person who created the note;
5. a model of the target users.

The first two have been used in various past and current geonote systems (such as [4,1]). However, indoor localization is still a relatively far less mature than outdoor localization based on GPS and used widely. In addition, we may need quite fine grained localization, since many exhibits might be in a single cabinet, especially for small ones, like spiders. We now explain the other three elements in relation to the way that the architecture operates. The last two

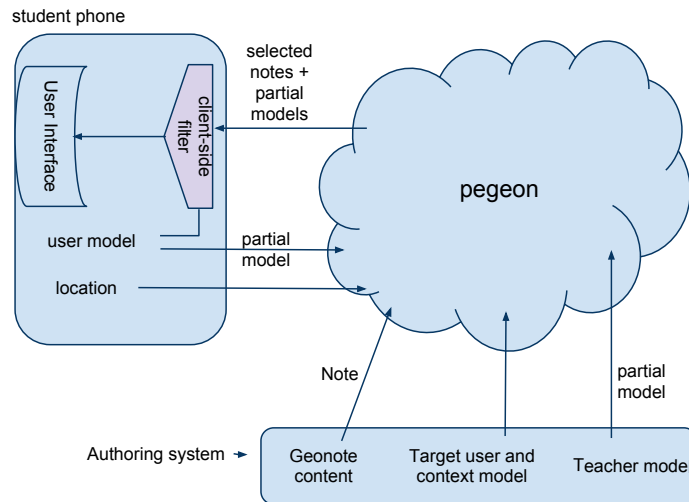


Fig. 1: Personalized GeoNote Architecture

elements are small or partial user models that act as metadata for the process of filtering.

The author interface must support creation of the actual content of the note. This may mean selecting images, writing text or recording audio. The author optionally defines the context model of the target user. At its most basic, this has the location associated with the note. More sophisticated context might include relationships between notes based on a prerequisite structure within the geonote system. It means that the teacher can easily create a series of notes, such as

- This case has a Sydney Funnel Web Spider. It is very dangerous. It likes to live in moist forests. It lives in lots of places around Sydney.
- This is a Black House Spider. It looks a lot like the funnel web spider you saw before. And people often confuse them. This one is not nearly as dangerous. But it is still a good idea to leave them alone.
- Look in this display for the Sydney Funnel Web and the Black House Spider. See how you can tell them apart.

The target user model is information about the person for whom the note is intended. In our scenario, all the teacher's notes are for a Grade 3 student, learning about spiders, as part of the Living Things syllabus component, The actual syllabus has several descriptors which are already used at the Australian Museum. For example, for the spiders, they include *Science and Technology* describing the *content* as *Living Things - LT ES1.3*. While some of this rather cryptic, it is used by teachers and so should be a natural way to tag learning resources and lesson plans. These can complement the more intuitive tags, such as *spider* or *Sydney Funnel Web*. Some notes have more specific target users.

For example, the teacher may provide notes with more pictures and less text for students with lower reading age. She may provide poems for students who have good reading skills and like poetry.

The last part of the authoring process is the definition of the author's partial user model. We show that the author's full model is stored on their own machine. They then define a *partial model* to be associated with the note. For our scenario, this could indicate that the author is a Grade 3 teacher in Sydney City. The teacher may choose to add their name and contact details (or not).

When the note is added to the Pegeon server, it carries these additional context and user models. This must store the notes and the associated context and user models. In a deployed system, this would need to deal with pragmatic issues such as security and privacy, and possibly a reputation system.

To this point, the description has been framed in terms of our scenario, with the teacher as the author of the geonote, at a conventional computer authoring interface, making use of maps provided by the Museum's education department. We also envisage a lighter weight interface that could readily operate on a phone. Then many people could be authors, including students, perhaps even Grade 3 students. A museum visitor might then elect to see only notes from teachers, or perhaps expert entomologists, or children. In addition, while the description has the full set of potential personalization elements, a set of defaults could support easy creation of notes, especially for the incremental cases, such as Alice adding one more note during the excursion, or David's creation of continuation notes.

We now consider the student phone system. When a student phone connects to the Pegeon server, it matches the partial model that they release to Pegeon as well as their location against the available notes. So, it does a first stage of filtering the notes at this stage. For example, if a student releases the part of the model indicating they are in Grade 3, studying spiders, just the notes matching these will be sent to their phone.

Each note still carries the context and user models that the author sent as metadata for the note. This means that client side personalization can be done on the phone. For example, the student's reading ability might be kept on the phone and used to filter the geonotes locally. For large geonotes, such as video, Pegeon might send just a link to the resource, rather than the actual content [11]. As we have already noted, indoor location is not well supported on mobile phones. In the figure, we show the phone as providing one important source of evidence which might be combined with other sources as we have done in previous work [2].

## 4 Discussion

Our extended scenario demonstrates two key ideas: the potential of geonotes in cultural heritage settings and the need for personalization to enable the visitor to find notes that are most relevant to them, at that right time, and context. In studies of the reasons that people visit museums, Falk [5] points to the potential importance of just this sort of personalization technology, accounting for the

user's motivations in the visit and the context. Our architecture is designed to support this level of personalization of geonotes. There are many quite challenging technical issues to be addressed. We now discuss these aspects briefly.

It will be critical to deal with ontological issues. How can we represent and use domain specific information? One approach is to make use of a museum ontology, collect evidence about user preferences and interests with respect to the specific museum ontology. This may work within a single museum but it then invites consideration of more general approaches that can operate across several different museums in the same general area, such as natural history. An alternative approach would be to use a generic common domain ontology. This does not currently exist (although it may in the future). A third, and potentially most appealing approach would make use of free text, as in existing widespread tagging systems. This makes for a challenging matching task if the system is to recognize which tags are similar; or it may simply mean that a system would fail to identify all relevant geonotes. A hybrid of all of these seems most promising. Notably, we have been exploring use of an ontology based on a formal school syllabus as this is already used in conjunction with materials at museums. It can address some parts of the ontological challenges. For other parts, free tags may be used.

Implementation of our architecture requires the following key elements: representation and management of the user models; management of the geonotes database and efficient management of multimedia resources. We plan to implement the system using our context and user modelling framework, Personis, for the central server's management of user models [2]. For the management of user models on the student's phone, we will use PersonisJ [8]. This will need to manage models of preferences, interests and knowledge, as well as contextual aspects, including spatial, temporal, social and possibly environmental details. For the management of geonotes, including arbitrary media, we are exploring approaches to content re-purposing so that we can deliver the note effectively to any mobile phone as well as other devices that could be useful in the museum context, such as an tablet or a tabletop.

The scenario and our description of the architecture have touched on the challenging and important issues of effective management of privacy, trust and reputation. The systems cannot succeed unless authors are able to confidently control the release of the partial user model they want to go with their notes and visitors can be confident about the partial models they release to the central server. The credentials of the author may need to be validated, especially where the visitors are children. Simpler approaches may work for informal geonotes but in other cases, such as those described in our scenario, the museum will either need to approve each geonote or, more efficiently, have a mechanism for identifying trusted authors.

The system relies heavily on effective user interfaces. There is considerable additional burden in authoring, compared with widespread geonotes. Even for the visitor reading notes, the system only operates effectively if they can easily establish their own user model. This might be done explicitly, in which case both

author and reader need support for tagging or using the available ontology. For the person reading the note, it may use established approaches to automatically build an implicit user model, based on their rating of some notes, perhaps using critiquing approaches [3]. It could also provide a set of standard responses, such as *You have to see this*. It may require an easy to use context editor, building on approaches previously described [10] to define the contextual conditions for delivering the geonote.

## Acknowledgments

This research was supported by Australian Research Council Projects DP0774532, DP0877665. We are grateful for the helpful comments from the reviewers.

## References

1. GD Abowd, CG Atkeson, J Hong, S Long, R Kooper, and M Pinkerton. Cyberguide: A mobile context-aware tour guide. *Wireless Networks*, 3(5):421–433, 1997.
2. M Assad, D J Carmichael, J Kay, and B Kummerfeld. PersonisAD: Distributed, active, scrutable model framework for context-aware services. In *Pervasive*, pages 55–72. Springer, 2007.
3. L Chen and P Pu. Interaction design guidelines on critiquing-based recommender systems. *User Modeling and User-Adapted Interaction*, 19(3):167–206, 2009.
4. F Espinoza, P Persson, A Sandin, H Nystrom, E Cacciatore, and M Bylund. Geonotes: Social and navigational aspects of location-based information systems. *UbiComp 2001*, pages 2–17, 2001.
5. JH Falk. *Identity and the Museum Visitor Experience*. Left Coast press, 2009.
6. M Fleck, M Frid, T Kindberg, E O'Brian-Strain, R Rajani, and M Spasojevic. Rememberer: A tool for capturing museum visits. *UbiComp 2002*, pages 379–385, 2002.
7. L Gannes. Foursquare turns 1 with half a million users. *Gigacom, Trusted Insights and Conversations on the Next Wave of Technology*, 2010.
8. S Gerber, M Fry, J Kay, B Kummerfeld, G Pink, and R Wasinger. PersonisJ: mobile, client-side user modelling. In *UMAP 2010*, pages 111–122, 2010.
9. J Hightower, S Consolvo, A LaMarca, I Smith, and J Hughes. Learning and recognizing the places we go. pages 159–176. 2005.
10. S Jbara, T Kuflik, P Soffer, and O Stock. Context Aware Communication Services in " Active Museums. In *SwSTE 2007*, pages 127–135, 2007.
11. J Kay. Vive la difference! individualised interaction with users. In C. Mellish, editor, *IJCAI*, pages 978 – 984. Morgan Kaufman, 1995.
12. T Kuflik, J Sheidin, S Jbara, D Goren-Bar, P Soffer, O Stock, and M Zancanaro. Supporting small groups in the museum by context-aware communication services. In *IUI '07*, pages 305–308, New York, NY, USA, 2007. ACM.
13. Board of Studies. *Science and Technology K-6 Syllabus and Support Document*. NSW Government, 1993.
14. J Pascoe. The stick-e note architecture: extending the interface beyond the user. In *IUI '97*, pages 261–264. ACM, 1997.

15. O Stock, M Zancanaro, P Busetta, C Callaway, A Krüger, M Kruppa, T Kuffik, E Not, and C Rocchi. Adaptive, intelligent presentation of information for the museum visitor in PEACH. *User Modeling and User-Adapted Interaction*, 17(3):257–304, 2007.
16. Y Wang, L Aroyo, N Stash, and L Rutledge. Interactive user modeling for personalized access to museum collections: The rijksmuseum case study. *User Modeling 2007*, pages 385–389, 2007.