Trail Records and Ampliative Learning

Research Report BBKCS-02-10,
School of Computer Science and Information Systems,
Birkbeck College, University of London, 2002

Mark Levene
School of Computer Science and Information Systems
Birkbeck University of London
London WC1E 7HX, U.K.
m.levene@dcs.bbk.ac.uk

and

Don Peterson
School of Mathematics, Science and Technology
Institute of Education University of London
London WC1H 0AL, U.K.
d.peterson@ioe.ac.uk

Abstract

One form of learning is the re-structuring and filtering of our experiences. We present a conceptual model of this type of learning, which we call ampliative learning, based on the notion of trails through an information ecology. The model addresses both the physical and virtual aspects of navigation facing the learner. We outline the design of a pervasive device, the experience recorder, that generates trail records associated with a particular experience, for example a visit to a museum, and supports both the learner and teacher in recalling and developing their trails. This leads us to some final observations on the pedagogical opportunities associated with this technology.

Ampliative Learning

Sometimes we learn by re-organising and filtering our experiences. Imagine that you are a teacher taking a group of students on a visit to an exhibition at a museum. During the visit, the students wander round and experience the exhibits in one order or another. The process of learning, however, does not stop there. After leaving the museum, the students engage in class discussion of what they have seen, they cross-reference items from books and from the internet, and they organise their thoughts into an essay. As learning takes place, new information is acquired, but also, crucially, existing information is re-structured, and so turns into useful and individual knowledge. A structure derived from the original tour through the exhibition, for example, will be progressively edited and rearranged as the student develops their essay. This process will be referred to here as ampliative learning, defined as the advance of knowledge through the structuring and re-structuring of one's original experiences.

Phenomena of this sort are widely recognised. In knowledge management and library science, for example, it is acknowledged that re-organising and filtering information may make it more useful for a given purpose. In developmental psychology it has been argued that the 're-representation' of existing knowledge can take us from 'basic competence' to improved performance (Karmiloff-Smith, 1992). A problem or puzzle can become much easier to solve when it is re-represented so as to make some things more salient and others less so (Peterson, 1994; 1996). And in the study of 'experiential learning' it has been argued that empirical processes of learning typically occur in a cycle which includes a 'reflective' process in which experiences are re-organised (Kolb, 1984).

The concern of the present paper is with the design of information technology to support ampliative learning after an experience such as a museum visit. We first offer a model of ampliative learning in terms of 'trails' through 'information ecologies'. We then present the idea of a 'trail record' for virtual and physical spaces. We then outline the design of an 'experience recorder' which creates trail records. Finally, we consider some possible strategies in pedagogy using this technology.

A Trail Model of Ampliative Learning

An environment in which we learn can be characterised as an information ecology, defined as 'a system of people, practices, values, and technologies' (Nardi and O'Day, 1999). Such ecologies include libraries, schools, museums, hospitals, the workplace, the home, and they may include physical space, cyberspace or a combination of the two. As we navigate through an information ecology, we enact a trail, which can be viewed as a connected sequence of items of information that were encountered during the navigation process. And as we learn, our trails are improved, extended and rearranged. In this sense, ampliative learning consists in trail refinement.

The notion of a trail was introduced by Vannevar Bush in describing his vision of a 'memex' machine (Bush, 1945; 1991). This would hold great quantities of information, and would allow the user to define paths through this store. As learning progressed, the user would edit and develop these trails, thus imposing new structures on the information held in the machine. For example, a researcher using the machine might define a trail through a series of books, papers and notes, and might develop this structure as his or her research proceeded. In this model, a trail is deliberately programmed into the machine by the user. A complementary idea is that a machine should record the trails which a person takes through an information ecology, presenting these as a trail record which can be stored, re-enacted or edited, thus supporting the cycle represented in Figure 1.

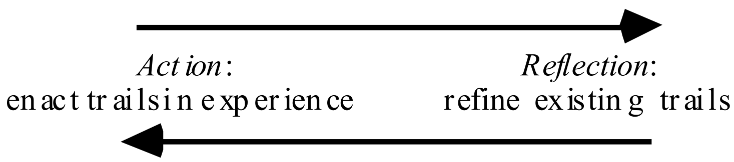

Figure 1: Cycle of Enactment and Refinement of Trails

The navigation of an information ecology can be understood as follows. A trail is more formally defined as sequence of places that were visited by the user within a navigation session such that the order of places within the trail is temporally coherent. A place occurs in the trail as many times as it was visited within the session; and a place may be physical, for example exhibits viewed in a museum, or virtual, for example web pages browsed. With each place we associate content, which gives meaning and attaches information to the place, and context, which describes how the user is situated in time and space when visiting the place. A trail record is a hypertextual trail providing an account of a user navigation session, be it physical or virtual. We assume that a collection of trail records can be browsed and edited using state of the art web/hypertext technology. Trail records provide us with a model of the users' actions and they can be viewed as a spatial/temporal account of their activities. It is possible to give trail records a logical basis, with reference to the appropriate spatial/temporal logic (Rescher and Urquhart, 1971) but as was shown by Levene and Loizou (1999) there are severe computational limitations to this approach.

We now turn to the issue of technological support for ampliative learning. Bush envisaged a machine which would not only store vast amounts of information (on microfilm), but would also allow the creation and refinement of trails through this 'maze of materials' (Bush, 1945). This, Bush thought, would provide a means of navigation which is faithful to the way the human mind actually works. That is: we do not naturally recall information by random indexing, but through the mediation of trails as defined above, and the 'memex' machine which he envisaged would support the specification and management of these trails.

These thoughts were published in 1945, before the proliferation of digital technologies and the World-Wide-Web (Berners-Lee, 1999). Computing devices are now becoming pervasive(Weiser, 1993): not simply tools which we use from time to time but part of the environment we interact with. Thus the information ecologies in which we navigate and learn are increasingly a combination of physical space and cyberspace. These developments, of course, pose a problem of information overload. Or to be more precise, they pose the navigation problem of moving in an effective and meaningful way though a vast ocean of available information. As Bush realised, this is not simply a problem of creating faster machines or creating more powerful indexing mechanisms such as those based on categorisation. Rather, it is a problem of supporting the management of the trails or associative clusters which the human mind naturally uses to mediate its relationship with its information ecologies. It is interesting to note that the infamous navigation problem of 'getting lost in hyperspace' (Nielsen, 1990) has been widely researched within the hypertext community (Levene and Loizou, 2002) and has several similarities with the issues of navigating within physical spaces (Darken and Peterson, 2002).

The facilities envisaged by Bush in his 'memex' machine have a directive character in which the user instructs the machine so as explicitly to specify and edit a trail. The scheme presented here has the further facility that the system automatically detects and records the trail which the user takes through a network and creates a record of this trail. This trail record can then be examined, manipulated and refined so as to support the process of ampliative learning. The scenario we are interested is, in broad terms, as follows. A user is navigating through a physical space and the experiences of his or her navigation session are recorded. Then the record of the physical navigation session is transported into virtual space as a hypertext. Finally, the user can navigate through this hypertext, which can be augmented with further content and hyperlinks within a virtual community of other users. In this way trail records can be seen to provide a fluid model of user action and ampliative learning, interfacing between the physical and virtual space dynamics.

An important feature of our proposed model is the provision for adaptive trail records, where the machine is given autonomy through intelligent algorithmics to learn users preferences from their actions over time and as a result provide more accurate and personalised access to the hypertext. Such adaptive records may result in augmentation of relevant content and hyperlinks to the original record and navigational assistance to the user trying to locate information within the record, as elaborated below. We find it significant that Bush revisited his original ideas regarding his memex hypertext machine and envisaged that the memex could 'learn from its experience' and 'refine its trails'. Oren (1991) calls this extended version the adaptive memex and discusses various adaptive technologies that could be used to realise Bush's vision(cf Nyce and Kahn, 1991).

Trail Records in Virtual Space

Assume a user is navigating through a web site. There are two points of view to consider: the user's view and the web site's view. Both the user and the web site can keep a record of the navigation session. Web browsers store history lists of the user's navigation, and web servers store log files of user requests to the server and, with the aid of 'cookies', which are stored on the user machine, web servers can identify the user on subsequent navigation sessions. Cookies and other forms of user identification by external entities raise various privacy and security considerations which we will not further discuss here. The 'correct' mode of operation depends on the application, for example, if we are navigating a commercial site the user may opt to be anonymous, while in an educational scenario it may be beneficial for the user to be identified to allow personalisation of the presented material.

Web data mining (Masand and Spiliopoulou, 2000) is a recent area of research which deals with the analysis of usage and the discovery of content on the web. In our context trail records can reveal patterns of user behaviour, and this knowledge can be used to optimise the interaction of the user with the web. The model elucidated in (Levene and Loizou, 2002) is a dynamic model, where the underlying hypertext (in this case the web), is viewed as a probabilistic automaton that accepts the trails traversed by the user. Ultimately the notion of adaptive trails is important as it allows the possibility of learning from the user's history in order to make inferences about present and future navigation sessions, and these will be used to improve the user's experience.

The adaptive framework is a point of departure from traditional computational systems, and gives greater autonomy to the system via machine learning algorithmics (Mitchell, 1997), for example based on a Bayesian standpoint (Lindley, 1985). The implication of this approach is that designers of adaptive systems do not need to anticipate every possible user scenario but rather build models of user behaviour which are flexible and can be manipulated according to individual user records. We believe that human-computer interaction should evolve in time rather than have rigid logical rules which are embedded in the machine and can only be changed by experts. This is not to say that expert knowledge is not imperative but rather that some of that knowledge is individual to the user, and may change and be adapted over time.

This will be especially important in an educational context. The individual student, with particular interests, will understand their own trails in an individual way. There is no one right way of conceptualising or representing a trail, and a system which is sensitive to the user's habits and preferences will be advantageous. The trail records which the system produces, therefore, should, over time, be adapted to the individual user. However, a degree of standardisation may also be desirable in contexts such as class discussion in which the records need to be compared and contrasted.

Trail Records for Physical Space

Now, we assume that the user is navigating through a physical space such as a museum exhibition. Again there are two points of view, the user's view and the view of the owner of the space. Users would like to keep a record of their experiences as a memory aid and as a means of extending, sharing and enabling better articulation of these experiences. In analogy to maintaining a diary, photo album or video recording, there should be provision for such records to be further edited, expanded and merged. The owner of the space would like to keep a record of users' activities as they were navigating through the space, in a similar sense to recording web logs for data-mining. Apart from gaining a better understanding of how people negotiate the space, user patterns that emerge from the mining process may help the owner to improve various aspects of the space such as access, layout and the provision of navigational aids.

Models of spatial cognition describe the representations of processes that relate to location based behaviour such as navigation from one place to another (Hartley and Burgess, 2002). Spatial behaviour involves place recognition, direction orientation and cognitive maps, which are mental representations of space involving topological information. Navigation can be enhanced by the use of various tools such as maps providing topological information, landmarks which are memorable places, trails in analogy to footprints through the space, and direction finding and orientation tools such as the compass (Darken and Peterson, 2002).

The trail record of a navigation session through a physical space is ultimately stored within a hypertext thus creating a connection between physical and virtual spaces; the hypertextual record can be further manipulated and enhanced with additional content and hyperlinks. One way to make the connection clear is to embed the hypertext on top of a map of the physical topology that was navigated, so that the hypertext provides a map of the portion of the physical space that was traversed. Using this technique the hypertext may become part of the wider web, when it links to related content, and receives links from records which are shared by a community having common interests or learning objectives. Adaptivity can also be built into physical trail records much the same way as when navigating a virtual space.

The loop can be closed when we revisit the physical space, and our new experiences augment the original record. With regard to the psychology of learning, it is important that this should happen. Ampliative learning, as we have defined it, is a re-structuring of experience, not a dislocated activity. The system described below is therefore designed to support a loop in which we enact a trail in actual experience, we refine and develop this afterwards, and the result then informs our return to actual experience.

The Experience Recorder

We now briefly describe a device, which we call the experience recorder, that allows the transporting into a (virutal) hypertext of navigation experiences within a physical space. To simplify the discussion we will continue to use the example of a person navigating through a museum exhibition. We will avoid implementation details in our description but stress that it is within the remit of current technology.

The experience recorder is a small, lightweight wearable computing device, the size of a badge, that has the following computing capabilities:

1) It can monitor the orientation of a user in the exhibition and embed the user's trail through the exhibition on top of the exhibit map, with temporal information relating to the times the user spent at various landmarks.

2) It can detect landmarks within the navigation session, for example the user viewing an exhibit. When such a landmark is detected the recorder can interact with it. The interaction may involve:

(i) Downloading an image of the exhibit, or a hyperlink to the image.

(ii) Downloading a hyperlink to content about the exhibit such as text, or, more generally hypermedia.

(iii) Taking a picture (or a video) of the user viewing the exhibit, or a hyperlink to such a picture.

3) It can record the user's voice when instructed to do so.

After a navigation session the experience recorder can transfer its log to a post-processor that converts this log into a hypertextual trail record embedded on a map of the physical space that was navigated. The trail record is then sent to the user, who can navigate through it on his or her browser and augment it with additional content and hyperlinks as any other hypertext. The trail record can thus be viewed as a web site that can be shared with other users. In a teaching situation trail records of different students can be discussed enabling them to share and refine their knowledge. As with any learning technology, the details of how best to use the experience recorder for educational purposes will need to be discovered in practice. However, we can envisage the form of its contribution, as outlined below.

Trail Record Pedagogy

The model presented above poses a clear pedagogical question: if learning is in part ampliative, and this consists in trail refinement, how can this be supported in teaching? More specifically: how can the proposed technology of the experience recorder be incorporated in the process of teaching? To illustrate, we return to our museum example. The students make some preparations before going to the museum. Once there they navigate through some of its exhibits. And later they discuss their experiences in class and write an essay. The proposed technology fits well into this larger scenario. In addition to classroom preparation for the visit, the students may be able to log on to the part of the museum's web site relevant to the visit. On the web site they see the layout of exhibits, and perhaps some standard recommended trails, represented in a form which is consistent with that of the trail record which they will be given on leaving the museum. On arrival, the students are given an experience recorder, as described in the previous section. On leaving, each student is given the web address of their personal trail record: a representation of their navigation of the museum's ecology, in the form, for example, of a web site. Back in class, discussion is facilitated by projecting several students' trail records on the wall simultaneously, providing a basis for discussion, explanation and debate. And when asked to write an essay, students use their trail records as a starting point, using computer tools to extend, refine and transform them into a finished document. In this way, the proposed technology makes actual and possible trails both explicit and manipulable, thus facilitating both the private and interactive aspects of ampliative learning.

In this sense, we envisage a trail pedagogy concerned with the effective and helpful use of trail records in teaching. The work of a lawyer, for example, involves trails through the complexities of a case, and in a team of lawyers, different members will have different but connected trails to build and follow as the case unfolds. In the teaching of law, therefore, it should be useful to represent and discuss these trails explicitly, and since their elements are increasingly available in electronic format, it would be possible to build a recorder which automatically generates trail records to support this sort of teaching. In this way, as in the previous example, projections of alternative trails generated by students in an exercise could be used to facilitate discussion, critique and explanation of how they have navigated among elements such as statutes, past cases, written communications, and their own notes. A doctor, likewise, navigates a complex ecology of procedures, locations, cases, regulations, machinery and online databases. A research scientist navigates an ecology of books, papers, libraries, online materials, and personal communications. A tourist navigates among historic sites, museums, guidebooks, monuments and other attractions. In some cases the function of an experience recorder would be auto-didactic, allowing the user to see, store, recall and ponder their own trails, and in others it would be instructional, allowing the teacher better to interact with the student in terms of their trails.

An experience recorder as envisaged here relies, of course, on electronic access to the elements of the relevant ecology. In the case of tourism, for example, this is barely possible at present: the relevant monuments or other attractions would not generally be visible to the recorder. However, the point for the future is that as information technology becomes more pervasive and ambient, the potential for such visibility will increase. This will naturally carry dangers as well as benefits: where there is an experience recorder, there could also be covert surveillance, and where there is explicit representation of trails, there could also be the attitude that one trail is right and others are wrong. The technology proposed here is not in itself inimical to privacy or diversity, but as with other such developments, its most important aspect will be its manner of use.

References

Berners-Lee, T. (1999). Weaving the Web. London: Orion Books.

Bush, V. (1945). As we may think. Atlantic Monthly, 176, 101-108.

Bush, V. (1991). Memex II. In Nyce, J.M. and Kahn, P. (eds). From Memex to Hypertext: Vannevar Bush and the Mind's Machine. Academic Press, New York, 165-184.

Darken, R.P. and Peterson, B. (2002). Spatial Orientation, Wayfinding, and Representation. In K. Stanney (ed), Handbook of Virtual Environments, Mahwah, NJ: Lawrence Erlbaum.

Hartley, T. and Burgess, N. (2002). Models of spatial cognition. In L. Nadel (ed), Encyclopaedia of Cognitive Science, London: Palgrave MacMillan.

Karmiloff-Smith, A (1992). Beyond Modularity: A Developmental Perspective on Cognitive Science. Cambridge, MA: MIT.

Kolb, D. (1984). Experiental Learning. New Jersey: Prentice Hall.

Levene, M. and Loizou, G. (1999). Navigation in Hypertext is easy only sometimes. SIAM Journal on Computing, 29, 728-760.

Levene, M. and Loizou, G. (2002). Web Interaction and the navigation problem in Hypertext. In A. Kent, J.G. Williams and C.M. Hall (eds), Encyclopedia of Microcomputers, New York: Marcel Dekker, 381-398.

Lindley D.V. (1985). Making Decisions. London: John Wiley.

Masand, B. and Spiliopoulou, M. (eds) (2000). Web Usage Analysis and User Profiling. Lecture Notes in Artificial Intelligence. Berlin: Springer-Verlag.

Mitchell, T. (1997). Machine Learning. New York: McGraw-Hill.

Nardi, B.A. and O'Day, V.L. (1999). Information Ecologies: Using Technology with Heart. Cambridge, MA: MIT.

Nielsen, J. (1990). Hypertext and Hypermedia. Boston MA: Academic Press.

Nyce, J.M. and Kahn, P. (eds) (1991). From Memex to Hypertext: Vannevar Bush and the Mind's Machine. San Diego, CA: Academic Press.

Oren, T. (1991). Memex: Getting back on the trail. In J.M. Nyce and P. Kahn (eds), From Memex to Hypertext: Vannevar Bush and the Mind's Machine. San Diego, CA: Academic Press, 319-338.

Peterson, D.M. (ed.) 1996: Forms of Representation: an interdisciplinary theme for cognitive science. Oxford: Intellect Books.

Peterson, D.M. 1994: Re-representation and Emergent Information in Three Cases of Problem Solving. In Dartnall, T.H. (ed.), Artificial Intelligence and Creativity. Dordrecht: Kluwer, 81-92.

Rescher, N. and Urquhart, A. (1971). Temporal Logic. New York: Springer-Verlag.

Weiser, M. (1993). Some computer science issues in ubiquitous computing. Communications of the ACM, 36, 75-84.