

## Trail Records and Navigational Learning

DON PETERSON

Institute of Education, University of London, UK

MARK LEVENE

Birkbeck College, University of London, UK

**ABSTRACT** An emerging wave of ‘ambient’ technologies has the potential to support learning in new and particular ways. In this paper we propose a ‘trail model’ of ‘navigational learning’ which links some particular learning needs to the potentialities of these technologies. In this context, we outline the design and use of an ‘experience recorder’, a technology to support learning in museums. In terms of policy for the e-society, these proposals are relevant to the need for personalised and individualised learning support.

### Introduction

As we enter the 21st century, the conditions of knowledge and learning have been changed irreversibly by the revolution in information and communications technology (ICT). The 1990s in particular saw an explosion in the use of the World Wide Web, and the emergence of a globalising knowledge economy in which the learner or enquirer has access, choice, flexibility and empowerment as never before. This is no Digital Utopia: we still have a divide in which many people are denied the necessary technologies and literacies; information overload and violation of privacy are increasing; and in school, homes and workplaces the facilities which have so far been created may be far from perfect. However the relation between the individual and the universe of information has changed permanently, as it also did with the invention of printing, and we need to leverage these conditions so as to ‘accentuate the positive’ and turn them to the good.

One such condition, attaching both to formal and informal learning, is the increased importance of navigation. As the mass of available information grows, so does the challenge of selecting and visiting what is relevant to our needs. We have greater choice in what we use, and equally we have greater responsibility for managing our own selections and pathways through these expanding oceans of information. The student who had access to a limited set of paper resources at least had the advantage that these may have been carefully selected, revised and tailored to his or her general needs. The same student given an open door onto cyberspace needs somehow to find and select what is relevant, suitable and useful in order to construct meaningful personal knowledge. And from formal education to the

most casual informal learning, this challenge of effective navigation magnifies as ICT proliferates.

In this paper we ask whether the child can alleviate the sins of the parent: whether new technologies have the potential to support something which has become more urgent, more critical due to the first wave of the ICT revolution. This something we call 'navigational learning', and the technologies we appeal to belong to the second or 'ambient' phase of the ICT revolution. We propose a 'trail model' of navigational learning, which links the needs of the learner to the potentialities of these technologies. We discuss 'trail records' in physical and in virtual space, and their use in an 'experience recorder' designed to support navigational learning in the context of a museum. And finally we discuss some issues of 'trail pedagogy' associated with such systems.

In terms of policy, these proposals are relevant to the increasing realisation by governments that the forces of globalisation require personalised and individualised learning support for citizens of the e-society. If a diverse, distributed and heterogeneous society is to function, then the individual needs of its members must be accommodated in e-learning as in other areas: social coherence, in this sense, requires individualised support. This point is recognised in the European Commission's 'Lisbon Strategy' proposed at its Lisbon Summit in 2000, in the associated L-Change Report (2003), and by the UK's Department for Education and Skills (2003). The point of the present paper in this context is that support for individual navigation in learning is an essential part of this general agenda, and that emerging ambient technologies have the potential to provide this.

### **Navigational Learning**

The general need for navigation is not new, and we begin with three scenarios that illustrate this aspect of learning independently of the recent rise of ICT. Imagine that you are a teacher organising a museum visit for a group of students. There has been some previous preparation and discussion, and the task to be completed after the visit is to write an essay on the links between ancient Greek and Roman sculpture. You all arrive at the museum, and the students wander through the exhibits, sometimes finding their own path, sometimes following a guidebook, sometimes discussing and making notes as they go. You leave the museum and the next day in class you hold a discussion in which students describe what they saw and how they interpret it. The students then take their notes, their memories, their discussions and their guidebooks, and set about developing their essays.

Imagine again that you are a solicitor and that you are asked to provide internal advice to your firm on contract law in medicine. This is part of your lifelong learning: you studied five aspects of contract law several years ago, and you now need to link this knowledge with new statutes and decisions concerning the medical sector. One of your colleagues has some knowledge of these links and provides useful pointers and references. After two weeks of work, you edit your materials together and construct a presentation for your firm. Two years later you are asked to revisit the same task including relevant updates and refinements.

Imagine again that you are a tourist visiting a new city and its historic sites. Before the visit you use literature, maps and websites to prepare yourself. During the visit, you navigate around its sites, collecting a miscellany of notes, diaries, photographs and memories. When you return home you may try to edit these into some sort of ordered record. This record may be useful if you wish to remind yourself of what you did and saw, share your experiences with others, or prepare to revisit the same city. And it may contain links to other records, for example from a museum visit which you made long ago when you were at school.

These examples of learning in different contexts show some interesting common features.

- They all involve *trails* of learning objects (see Bush, 1945). That is: the museum exhibits, the legal statutes, the tourist locations, and so on, are not encountered in isolation, but as elements in longer paths.
- Some of the time we are engaged in *enacting* these trails: as we pass through the museum, undertake our course in contract law, or visit another site. We may be following a prescribed path from learning object to learning object, we may be creating something new and exploratory, or we may be doing a bit of both, but in any case we are ‘on the ground’ engaging in a series of experiences.
- Some of the time we are engaged in *editing* our trails: as we plan them beforehand and refine and extend them afterwards, reflecting, reorganising and making and deleting connections.

Thus, information becomes personal knowledge through a process of construction, and this process involves a cycle in which we enact, edit, re-enact and re-edit our trails. We will call this cycle ‘navigational learning’, using the general formulation:

navigational learning = enactment + editing of trails

The students in our scenarios create trails in the museum (enactment) and afterwards they discuss, refine and extend these (editing). The solicitor in our scenarios engages in navigational learning in which she undertakes a course in contract law (enactment), returns to this later, refines it and links it to new statutes (editing) with the collaboration of a colleague, presents a worked-out structure to her firm (enactment), and two years later returns and begins anew. The tourist in our scenarios shows a similar cyclical pattern: he edits together trails which are preliminary plans, he enacts, adapts or replaces these ‘on the ground’, shares records of the enacted trails with others when he returns home, and uses these records as a starting point when he prepares to return again to the same region.

These phenomena find various reference points in the theory of learning. The general position of constructivism emphasises that we build personal knowledge and meaning through action and interaction (rather than just by passively receiving transmitted information). Kolb (1984), in his account of ‘experiential learning’, emphasises the cyclical relation between experience (enactment) and reflection (editing) on what has been or will be experienced. Dewey (1916) emphasises that reflection (editing) supports the continuity of action and experience (enactment). And further back, Aristotle emphasises the reciprocal relation between enactment and knowledge when he says ‘what we need to learn before doing, we learn by doing; for example, we become builders by building, and lyre-players by playing the lyre’ (2000, II/1). In these terms, the general position taken here might be referred to as ‘navigational constructivism’.

### Trails and Computers

What is new, however, is the pressure on navigation created by the ICT revolution. We need now, more than ever, to find support for this aspect of learning as we try to make good use of ever larger oceans of available information. And we need to ask how we can use machines to help us with a challenge itself created by machines. Some key ideas here are provided by Bush (1945), in a paper written just before the first modern computer was built. Here Bush introduces the notions of a ‘trail’ and an imaginary machine, the ‘memex’. He considers the case of the scientist searching through libraries, forming trails of books, papers and

illustrations, improving these trails through recording, editing, and sharing them with others. And he suggests that a machine, using levers, microfilm and other pre-ICT technology, might help the user in recording and working with these trails. This early vision of how to deal with the challenge of navigation through a 'maze of materials' is distinctive in two ways: it treats trails as 'first class' learning objects along with the papers and books which compose them, and it treats them as entitled to their own tailored forms of mechanical support. Since that time the challenge of navigation has magnified, and is encountered by a wider range of learners in a wider range of learning contexts, and so we need to ask Bush's question again. As we would put it: how can ICT support navigational learning?

In the beginning the computer was a tool: a device which performed calculations, stored data, and with the advent of networking, provided access to information over great distances. This tool took over clerical functions in banking, control functions in power stations, and it even learned to play chess. In the form of the PC it became widespread in schools, homes and workplaces; and supported by the Internet, it has been making what used to be inaccessible and distant accessible and near at hand. However, the future story of ICT will not simply be a matter of the further proliferation and miniaturisation of devices with the same functions. A new set of ideas and technologies exists in research laboratories around the world, and is starting at the time of writing to enter the marketplace. This is the ambient phase of the ICT revolution, in which the computer gives us environments or ecologies (Nardi & O'Day, 1999) in which we learn, in collaboration with computer systems which surround us, and which adapt to our needs. This involves 'ubiquitous' hardware (Weiser, 1993): devices which are embedded in our furniture, public spaces, vehicles and clothing, and which may be aware of our presence. It involves 'semantic' software (Berners-Lee, 1999): ICT systems which filter, interpret and structure so as to participate in interpretation and the creation of meaning. It involves 'adaptive systems' which learn about an individual's habits and needs, and mould their behaviour accordingly, for example producing 'personalised services'. And it involves 'intelligent agents' which work for us to collect and transmit information, which work without our supervision, and which communicate and collaborate with each other [1]. None of this will be magical or perfect, and it will no doubt create problems as well as benefits. But neither is it science fiction: some of it has already entered e-commerce, more of it is ready for commercial exploitation, and we can expect these technologies to become commonplace during the first two decades of the 21st century.

This means that it is becoming insufficient to ask 'What can the computer, or ICT, do for learning?' Yes, the computer is a tool, but it is also an ecology. It is a device on the desktop or in the hand, but it is also an adaptive and potentially supportive environment. It is an object that we program, but it is also a personalised service and a symbiotic collaborator. This new paradigm does not eliminate the old one, but it supplements it significantly, and it extends the potential of the computer to support learning. In particular, it suggests forms of learning support which facilitate the recording and editing of our trails.

### **Support for Navigational Learning**

The learning objects which we encounter may be either physical or virtual (electronic): we might pass from a paper guidebook to a museum exhibit to a web page accessed from within the museum to a paper notepad, and so on, creating a coherent learning experience from all of these. A trail is defined as sequence of *nodes* that are visited by the user within a navigation session. With each node we associate *content*, which gives meaning and attaches information to the node, and *context*, which describes how the user is situated in time and space when visiting the node. 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 & Urquhart, 1971) but as was shown by Levene and Loizou (1999) there are severe computational limitations to this approach. 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 & Loizou, 2002) and has several similarities with the issues of navigating within physical spaces (Darken & Peterson, 2002).

Bearing in mind our definition of navigational learning as the enactment and editing of trails, we can state in general terms what is required of technology which supports this. Essentially, we need such technology to create trail records and to facilitate the editing of these records so that they can be examined, shared, manipulated, refined, stored and connected with one another. This suggests the general formulation:

support for navigational learning = support for creation + editing of trails

That is: we want the technology to watch us as we go, and to present us with a record that we can manipulate afterwards, store, and use again. This, of course, leaves plenty of room for detail.

One possibility is that the system provides guidance as to which trails may be of use to the learner. To do this it will need to know something of the learner and something of the learning objects in question: and, comparing the two, it may assess how relevant, appropriate or acceptable a particular learning object is to learners with particular backgrounds, preferences and learning styles. Research on this type of system can be found, for example, in the SeLeNe project funded by the European Union (Poulovassilis, 2003).

Another possibility is that the system 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 in 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 own trails. Oren (1991, p. 320) calls this extended version the 'adaptive memex' and discusses various adaptive technologies that could be used to realise Bush's vision (see Nyce & Kahn, 1991).

Another possibility is that the system allows us to conceptualise and represent our trails in different ways, giving us alternative views of our navigational activity. What for one learner is an alternation between European and American exhibits might be understood by another learner as a chronological progression, irrespective of geographical origin. Different goals and learning tasks require different representations, and this applies to the trails we enact and edit just as it does to our general ability to reason and solve problems (Peterson, 1996).

The scenario which we now focus on 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 navigational learning, interfacing between the dynamics of physical and virtual space.

### **Trail Records for 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's 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 & 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 by 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, so that these can be used to improve the user's learning experience.

This adaptive framework is a point of departure from traditional computational systems, and gives greater autonomy to the system via machine learning algorithmics (Mitchell, 1997), based for example on Bayesian principles (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 that 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 that 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 trail 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 & 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 & Peterson, 2002).

The trail record of a navigation session through a physical space is ultimately stored within a hypertext such as a web page, 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 that are shared by a community having common interests or learning objectives. Adaptivity can also be built into physical trail records in 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. The editing component of navigational 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 possible device, which we call the *experience recorder*, which allows the transporting into a (virtual) hypertext of navigation experiences within a physical space. The requirements for the experience recorder are, in general terms, to support navigational learning in a pervasive and non-intrusive manner. 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. (Our aim is to present a conceptual and computational framework for navigational learning, and therefore the detailed specification for a given application is out of the scope of this paper.)

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

- 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.
- It can detect landmarks within the navigation session, for example the user is 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.
- It can record the user's voice when instructed to do so.

The experience recorder works together with an ambient system, embedded in the architecture of the museum. 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 and can be shared and discussed with others. From the user's point of view, this could involve registering online before the first visit to the museum, wearing an unobtrusive badge-like device during the visit, receiving by email the address or URL of their personal trail record web page, making use of this as described above, and perhaps returning to it years later. We envisage that this may increase not only the learning potential of the visit, but also the individual's sense of owning their learning experience.

### **Trail Record Pedagogy**

In order to facilitate the use of an experience recorder, it will be important to guide students so as to situate its employment within the larger cycle of navigational learning. In our museum scenario, for 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 that is consistent with that of the trail record that they will be given on leaving the museum. As described above, the students are given an experience recorder on arrival, and the web address of their personal trail record on leaving: a representation of their navigation of the museum's ecology. 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 navigational learning.

The work of a lawyer, to take another example, involves trails through the complexities of cases, and in a team of lawyers, different members will have different but connected trails to build and follow as the case unfolds. Thus in the teaching of law, 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, emails, and their own notes.

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. The point for the future is that as ICT becomes more pervasive and ambient, the potential for such visibility will increase. This will naturally carry dangers as well as benefits:

the possibilities of covert surveillance, confinement through the limitations of the technologies and so on. In any case, as the e-society develops, the computer-as-environment will change both the conditions of learning and the possibilities of learning-support, just as the computer-as-tool has already done: and what we propose here is a framework for making good use some of these changes.

Correspondence: Don Peterson, School of Mathematics, Science and Technology, Institute of Education, University of London, London, WC1H 0AL, UK. Email: D.Peterson@ioe.ac.uk

## Note

- [1] Web searches on terms such as ‘ubiquitous computing’, ‘semantic web’, ‘adaptive systems’ and ‘intelligent agents’ should produce material on these developments, and adding the terms ‘conference’ and ‘abstract’ should elicit material which is more up to date than what is currently in print.

## References

- ARISTOTLE (2000) *Nicomachean Ethics*, Cambridge, Cambridge University Press.
- 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: J.M NYCE & P. KAHN (eds) *From Memex to Hypertext: Vannevar Bush and the mind’s machine*, New York, Academic Press.
- DARKEN, R.P. & PETERSON, B. (2002) Spatial orientation, wayfinding, and representation. In: K. STANNEY (ed.) *Handbook of Virtual Environments*, Mahwah, NJ, Lawrence Erlbaum.
- DEPARTMENT FOR EDUCATION AND SKILLS (2003) *Towards a Unified e-Learning Strategy*, Nottingham, DfES Publications.
- DEWEY, J. (1916) *Essays in Experimental Logic*, Chicago, IL, University of Chicago.
- HARTLEY, T. & BURGESS, N. (2002) Models of spatial cognition. In: L. NADEL (ed.) *Encyclopaedia of Cognitive Science*, London, Palgrave MacMillan.
- KOLB, D. (1984) *Experimental Learning*, New Jersey, Prentice Hall.
- L-CHANGE (2003) Change in European education and training systems related to information society technologies (IST). Yearly report 2002/2003. Available at: <http://www.education-observatories.net/lchange/>.
- LEVENE, M. & LOIZOU, G. (1999) Navigation in Hypertext is easy only sometimes, *SIAM Journal on Computing*, 29, 728–760.
- LEVENE, M. & LOIZOU, G. (2002) Web interaction and the navigation problem in hypertext. In: A. KENT, J.G. WILLIAMS & C.M. HALL (eds) *Encyclopedia of Microcomputers*, New York, Marcel Dekker.
- LINDLEY D.V. (1985) *Making Decisions*, London, John Wiley.
- MASAND, B. & SPILIOPOULOU, M. (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. & 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. & KAHN, P. (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 & P. KAHN (eds) *From Memex to Hypertext: Vannevar Bush and the mind's machine*, San Diego, CA, Academic Press.
- PETERSON, D.M. (1996) *Forms of Representation: an interdisciplinary theme for cognitive science*, Oxford, Intellect Books.
- POULOVASSILIS, A. (2003) SeLeNe: self e-learning networks. Available at: <http://www.dcs.bbk.ac.uk/~ap/projects/selene/>.
- RESCHER, N. & 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.