A Diary Study of Information Capture in Working Life

Barry A. T. Brown, Abigail J. Sellen, Kenton P. O'Hara

Appliance Design Group
Hewlett-Packard Labs
Filton Road, Stoke Gifford
Bristol BS34 8QZ, UK
+44 117 922 9520

Barbro@hplabs.hp.com, Abisel@hplabs.hp.com , Kenoha@hplabs.hp.com

ABSTRACT
Despite the increasing number of new devices entering the market allowing the capture or recording of information (whether it be marks on paper, scene, sound or moving images), there has been little study of when and why people want to do these kinds of activities. In an effort to systematically explore design requirements for new kinds of information capture devices, we devised a diary study of 22 individuals in a range of different jobs. The data were used to construct a taxonomy as a framework for design and analysis. Design implications are drawn from the framework and applied to the design of digital cameras and hand held scanners.

KEYWORDS
Information capture, appliances, digital cameras, voice recorders, scanners, diary study, PDAs, document use

INTRODUCTION
Advances in technology are increasingly enabling devices which allow for the capture of a wide variety of information or media, whether it be for capture of paper documents, scenes, sound, or moving images. The boom in digital cameras is perhaps the most salient example, but we are also seeing gadgets hit the market in the form of voice recorders, digital video cameras and document scanners. Not only are such devices becoming more ubiquitous, they are also becoming more personal: they are becoming smaller, more portable, and often more specialised for particular kinds of information capture. In addition, we are seeing more diverse and sophisticated software for delivering, manipulating, storing and viewing captured information. These capabilities, coupled with huge advances in computer storage capacity mean that new technology is greatly enhancing our ability to record and keep images of everyday life events.

These advances unleash great potential for technological innovation. For example, cameras or voice recording capabilities can be easily added to existing devices such as desktop computers, laptops, and palmtops. Cheaper, better capture technologies are also leading to new classes of device for the consumer market (e.g., the Sony Voice Balloon, the CoolPix camera) as well as the business market (e.g., the Crosspad, the Cpen). Since these are new classes of device, it is often not immediately obvious how these products might eventually be used by consumers. IT companies face important decisions about how to design such devices for optimal use, and how to market them.

In our own research group, we were confronted with these very issues in relation to a new kind of capture device, a handheld document scanner known as the "Capshare 910" (Figure 1). Marketing research had resulted in good positioning for the product in the mobile professional market. We wanted to understand whether other markets, and whether other kinds of activities might represent future opportunities for this kind of technology. We were also interested in changes to the hardware and software which could result in new types of Capshare type devices.

Fig 1. The Hewlett-Packard Capshare 910

To do this, we needed to first understand the range of situations in which people want to capture paper documents. We also wanted to understand the subsequent ways in which they then want to use that information. Another related question was whether or not document capture could be enhanced by considering extending the kinds of information that could be captured (such as adding

sound or scene capture). We felt that without a deeper understanding of these issues, we could not systematically map out the space of design possibilities.

The Literature

Unfortunately, the existing literature is not particularly illuminating on this topic. As it stands, it has little to say about the range of ways in which people capture information or the reasons why they do it. Rather, research has tended to focus on the design and use of specific systems which make use of capture technologies. So, for example, there are descriptions of systems based on camera input designed to capture paper documents on the desktop [e.g., 6]. There are other systems designed to capture office conversations and meeting events by using combinations of audio, video and scribble information [e.g., 2, 5, 11]. Other research reports have covered the use of active badges to automatically capture activities such as people's movements within a building [4]. In addition there have been many papers written over the years describing systems which make interesting use of the capture of moving images (including hand gestures) and still images or scene capture [e.g., 3].

From a design perspective such as ours, information capture has not been recognised as a topic in its own right. Perhaps because of this we found little which could offer guidance on what sorts of devices people might find useful. The literature as it stands does not discuss the circumstances under which people might *want* to capture information if they only had the means. Also, it does not give clues as to what kinds of captured media might be usefully combined (for example audio plus scene), and it does not point to the kind of software, services or supporting infrastructure that might be valuable to help people deal with captured information.

In response to this, we designed a study to analyse people's information capture activities. The paper starts by discussing the study, and the *taxonomy of capture* which we developed using the results. This taxonomy unpacks the diversity of capture activities leading to a number of design implications. In particular, we focus on implications for the hand-held scanner we have described, but we also extend the findings to other kinds of potential capture devices.

Our Approach

Since Capshare was marketed as a business product, one of our primary interests was to understand the range of capture activities that people do in very different kinds of jobs. The focus in this initial research was therefore on understanding information capture in workplace activities. We had no specific vertical market in mind, wanting instead a diverse cross-section of people.

These considerations led us to choose a diary method with a mixture of subjects from a range of occupations. Diary methods are a popular data collection technique in sociology, but are still relatively rare in technology studies.

This is a method we have used in the past with some success in the areas of reading [1], paper use [9], and the research behaviour of library users [7].

Finally, we wanted to design a study which would not only provide data on people's paper document activities (to inform future directions for Capshare) but also to look at information capture activities more generally. Thus, for half of the subjects in the study we focussed on their paper document capture activities; for the other half we were interested in any kind of information capture they carried out or wished to carry out.

METHOD

Choice of Subjects

A mix of subjects from both within and outside our organisation was recruited – 13 from within HP, and 9 from outside HP. All were required to be PC and email users.

Eleven of the HP subjects formed the multimedia capture group (Multimedia Grp). For this group we were interested in all of their information capture activities (not just documents). To recruit these participants, an email was sent out to the Hewlett-Packard Bristol site of 2,000 non-research staff. The subjects were then chosen on the basis of their interest and occupation. Here the aim was simply to recruit a diverse mix of both professional and administrative staff. The resulting sample included a marketing manager, a sales executive, an administrative assistant, a security manager and a logistics specialist.

Eleven subjects (2 from HP and 9 external) were recruited for the paper document capture group (Paper Grp.) This was done using a list of occupations which were chosen as highly "information intensive". In particular, because of our interest in hand scanners, we chose people whose jobs were highly reliant on the collection and capture of paper documents. The resulting pool of 11 subjects in the Paper Grp. covered a wide range of occupations including a lawyer, journalist, stockbroker, office administrator, financial administrator and teacher.

While 22 subjects constitute a relatively small sample size, it is appropriate for a diary study in which each person contributes in-depth data over the course of several days. Moreover, our aim was to broaden our understanding of what information capture means rather than to produce a set of data which is statistically generalisable. This mandated the in-depth study of a wide group of different occupations.

Procedure

This study used a diary methodology which was a modification of previous methods we have used. Instead of asking subjects to keep written notes on their activities, we asked them to take photographs of the events we were interested in. We equipped subjects with digital cameras to use over the course of 7 consecutive days (covering on

average 5 working days and 2 days at home). Subjects in both groups were asked to use the camera whenever during the course of each day they felt the need to "capture" some information either at work or at home. It was emphasised that they should use the camera as a *diary tool* rather than as a conventional camera. They were told to take a picture whenever they *actually* captured some information in the course of their day, or whenever they would have *liked* to have captured information but did not have the means.

The procedure for both the Multimedia Grp and the Paper Grp was the same with one exception. Subjects in the Multimedia Grp were told we were interested in opportunities to capture any kind of information they came across – be it spoken or ambient sound, document-based information (paper or electronic), moving image, or scenes. Subjects in the Paper Grp however were told that we were only interested in information captured from paper documents. These paper documents could take on various forms such as books, post-it notes, magazines, newspapers, as well as more conventional business documents.

For both groups, the pictures themselves were used later as illustrations and as memory joggers in semi-structured interviews intended to unpack the context surrounding each capture event. Subjects were interviewed three times over the week, and asked in detail about each photograph they had taken. For each photograph, they were asked a number of different questions including:

- What did they and how ideally would they have captured the information?

- How did they or how would they have used the information they captured?

- Did they share or did they want to share the captured information?

- Did they or would they have wanted to keep the captured information?

Each of these interviews lasted anywhere from thirty minutes to an hour. The interviews were tape recorded and fully transcribed, producing a large corpus of information. The *methodology we used has similarities with photo-elicitation*, a technique used in anthropology and sociology [8], with the exception that the study participants – rather than the researcher – actually takes the photographs.

RESULTS AND DISCUSSION

The study generated data on 381 "capture opportunities": 219 photographs from the Multimedia Grp and 162 from the Paper Grp. A capture opportunity is defined as an occasion on which they used the camera to record either an actual capture event or a potential capture event.

What The Multimedia Group Captured

As we hoped, the capture opportunities in the Multimedia Grp reflected a wide range of information media. As Figure 2 shows, subjects in this group indicated that they wanted

to capture moving images, audio (both voice and ambient sound), visual scenes, and information from a variety of different kinds of paper documents (containing for example, to-do items, hand-written notes, financial data, etc). There were also some kinds of information which were harder to classify. Thus the "Other" category includes opportunities for capture of electronic information (such as Web pages), projected information (such as slides), and other events where the "type" of information is not easy to classify (such as the desire to capture information off a car license plate).
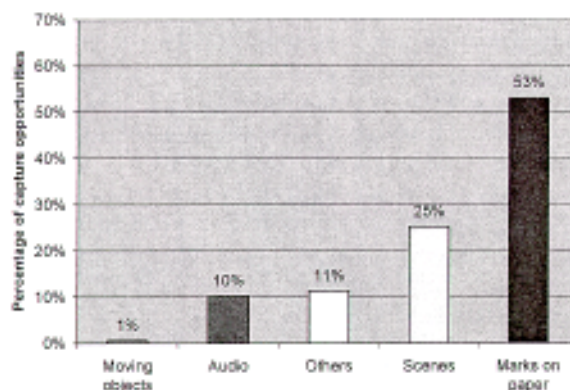

Fig 2. Kinds of information "captured" by the Multimedia Grp

This confirmed to us that subjects were able to think of the camera not necessarily as simply a camera but also as a recording tool. So for example, one subject took a picture of a colleague and told us that he had wanted to capture what that person had said; another took a picture of a meeting and told us it was a video snippet she had wanted to capture.

Perhaps because most of the diary data came from work situations, marks on paper constituted the most frequent kind of information subjects said they wanted to capture. Examples included capturing handwritten notes and photocopying paper documents. However, also prominent amongst these data were examples of subjects using the camera *as a camera* to capture scenes. This perhaps reflects a bias in use of the camera as a diary tool in this methodology. Nonetheless, what was interesting was the ways in which subjects found scene capture to be useful in different work situations. For example, one participant (we will call Phil) worked as the site security manager. Due to a large number of thefts from the site, Phil was attempting to have a security camera fitted, however he met with considerable resistance from staff who had hesitations about being "spied upon". To address these concerns Phil climbed up on a ladder and took a photograph from where he wanted to have the security camera fitted. Phil explained that he then not only used this photo in a report to management, but also showed it amongst staff to alleviate

their fears (Figure 3). As he put it when interviewed: "people can say they know, but when they see the photo they say ahh! Now I see!".

Fig 3. View from a potential security camera

In addition to this example, there were many others in which scene capture provided the most effective medium by which information could be captured and used. Photos were not only shared on the camera itself, but were also captured for inclusion in documents and for emailing.

Sound capture (mainly of voice) represented another frequent kind of information capture opportunity reported. Some of the subjects stood out as sound capture enthusiasts. For example, one subject (Steve), despite worries about ethics, discussed several occasions where he would like to have recorded ongoing telephone conversations. Steve worked on technical support for medical knowledge management systems. Throughout his day he often had long unexpected calls from clients asking for help with software problems. Steve would need to calm the client and ask the right questions while taking detailed notes. By tape recording his phone conversations, he felt he might be able to better concentrate on discussing the problem with the client, reassured that he could go over the tape recorded call and take notes later.

This analysis implies that there is indeed potential for devices in the workplace that capture both sound and scene. For instance, it suggests that there may be an important new market for digital cameras designed as workplace tools rather than as consumer devices. Perhaps not surprisingly, the analysis also points to the overriding importance of devices which can effectively capture paper-based information.

What The Paper Group Captured

An analysis of what the Paper Grp captured gives us further insight into the different kinds of paper documents that people extract information from (Figure 4).
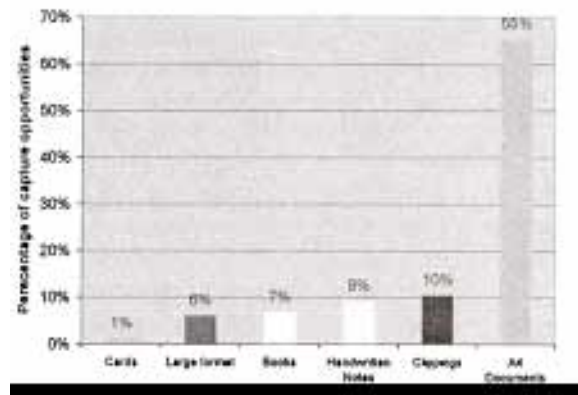

Fig 4. Kinds of paper-based information "captured" by the Paper Grp.

The documents took many different forms, from more conventional business documents (A4), to newspaper and magazine clippings, to books, large format documents (mainly flipcharts), and cards. By far the most common to be photographed, however, was regular letter-sized (A4) paper documents. Interestingly, over 62% of these were single sheets of paper, and less than 6% were documents of more than 3 pages. The second most frequent kind of paper document our subjects wanted to capture was that of clippings. These were called clippings as they represented occasions when subjects said they wanted selective pieces of text or graphics from newspapers, magazines, and other documents. Almost as frequent was the desire to capture handwritten notes and scribbles, books, and also large format documents such as flipcharts.

This kind of breakdown has implications for the kinds of document capture devices one might provide for users. These data suggest that although A4 documents are the most important type of captured document, 35% of documents that are captured are not A4. This suggests that a document capture device should have the ability to capture as wide range of different media as possible, or that one needs to think in terms of specialised devices.

Capture Taxonomy

More interesting than what people captured or wanted to capture, is *why* or *for what purpose* the capture activity took place. Capture of information is goal-oriented. That is, we do not capture information for its own sake (or at least not in most cases), but rather we do it to use that information in some other way. A key aim of the study was to elaborate on these goals since, as we will show, this significantly impacts the nature of the devices designed to support these goals.

To highlight these differences we sought to construct a taxonomy of the different capture opportunities we had collected. We did not want to develop an independently verifiable taxonomy, or a categorisation scheme which was

necessarily exhaustive and comprehensive. Further, we were not constructing this in order to test any hypotheses or conduct any experimental tests. Rather we sought to come up with a descriptive framework: one which would help us (and others) to understand the diversity and range of capture activities that people want to carry out; and perhaps more importantly, one which would make useful distinctions amongst activities for the purpose of design of new technologies.

The categories were thus derived as a result of extensive discussion of the data to draw out similarities and differences between the different capture opportunities we had collected. The capture situations were categorised according to the reasons behind the capture as expressed by the subjects in each of the events in our corpus. This resulted in a ten category taxonomy which exhaustively categorised all the capture events. For some of the capture events there were multiple reasons for information capture, so the capture categories are not mutually exclusive. The categories of capture can be briefly described as follows:

1. Capture to discuss
Items which have been captured in order to have an interactive, synchronous conversation around the captured information or document. In this case the captured item is a part of a spoken conversation. This can be done with those physically co-located, or those who are distant, e.g. over the phone

2. Capture to distribute
Items which have been captured to be distributed or sent to someone. Typical methods include hand delivery, fax, email. This is different from the first category in that the document can be sent without the need to have a conversation around it. The intended receiver(s) may be remote or co-located.

3. Capture to post in a common information space
Items which are captured to put them into a common space, such as a shared desk, a notice board, a shared whiteboard or a fridge door. These items are often positioned in prominent places so as to allow 'incidental viewing' – noticing the item without specifically looking for it.

4. Capture to archive
Items which are captured for long or medium term storage. They are typically items which are captured with no specific short term purpose in mind, but are kept because there is a feeling that they might be useful 'just in case'. Materials produced in meetings often fall into this category.

5. Capture to collect and collate
Items which are captured to add to a collection of similar items. This might include collecting papers for a personal library, clippings for a magazine, interesting URLs, and so on. The important aspect here is that the collection unites objects with some sort of common theme, and it is the collection as a whole which is of worth.

6. Capture to read and reflect
Items which are captured to be read or reflected upon at a later time when more convenient. This is often the reason for "collecting" documents so they can be taken back to base and read later.

7. Capture for task management
Items which are captured in order to help remember things that you have to do, or have to be organised in the future. Examples include task lists, post it notes and so on. Many recent inventions such as PDAs, attempt to cater for this use.

8. Capture to refer to
Items which are deliberately captured to be referred to in some later activity. This includes situations such as capturing a phone number for later use, taking notes during a presentation to refer to when composing a new document, or capturing an advertisement to order a product. This does not include situations where information is captured so that it can be incorporated into a document verbatim (see next category).

9. Capture to re-use
Items which are captured in order that they can be re-used in the production of something else (usually a document). Examples include capturing a slide for use in composing a new presentation or capturing text or figures to incorporate into a new document. The emphasis here is on the *verbatim* reproduction of what was captured.

10. Capture to a living document
Items which are captured and used as "living documents", with the document being written on, annotated and edited through the day. "Living documents" are typically modified in an ongoing way, and act as structured information repositories. A project plan is a good example of such a document.

Fig 5. Frequency of capture by category (Multimedia Grp.)

Fig 6. Frequency of capture by category (Paper Grp.)

Using this taxonomy we were then able to analyse the frequency with which subjects in both groups carried out these different kinds of capture activities. Figures 5 and 6 show differences in the most frequent kinds of capture activities the two groups carried out. The most common form of capture for the Multimedia Grp was capturing documents, scenes and even audio snippets for use as a reference in a later activity (*"capture to refer to"*). The Paper Grp, on the other hand, mainly wanted to capture

documents for the purpose of sending or distributing the information to others (*"capture to distribute"*).

Rather than attempting to generalise from these somewhat small and diverse samples, the important point to take from both the taxonomy and the graphs is to recognise the very different kinds of goals people have in mind when they capture information. Further, the most frequent kind of capture activity that any one person does is likely to be highly dependent on who that person is and the kind of job that they do. Some subjects, for example, acted as something of an information hub – they would gather information from diverse sources and redistribute it to others. Accordingly, these subjects mainly *captured to distribute*. An example of this was the administrative assistant we studied. Other subjects tended to collect information for future reference and use when creating new documents – *capture to collect, capture to refer to* and *capture to re-use*. The journalist we studied was an example of this in that he often collected newspaper clippings to use when he was writing new articles.

APPLICATION OF THE TAXONOMY

We turn now to why we believe this kind of analytic framework is both useful and interesting. The first and perhaps most obvious point is that such a taxonomy is didactic – it provides a framework and a language with which to talk about information capture. In doing so, one can better understand how any particular capture event, or the capture activities of any one person fit within a broader picture.

The second point, however, is more practical. Each of the capture categories has implications for designing technology which would support those activities. To illustrate this, we will consider four of these categories and the particular implications these have for two types of capture technologies - Capshare, and digital cameras.

Capture to discuss

When people used captured information for the purpose of discussion, being able to jointly view the captured information was key to the activity. Thus, using currently available technologies, people in meetings displayed information by means of projectors, large display screens, and more old-fashioned technologies such as paper.

Of interest was that subjects also used the LCD screen on the back of the camera for the purpose of discussing captured information with others. For example, one subject used a photograph of a faulty power supply to convince a supplier that it was incorrectly fitted. The picture became an important part of the discussion, and the ability to look at it together enabled this to happen.

Once we were sensitised to these kinds of activities, we began to realise the extent to which people were using the Capshare for discussion around shared images. In doing so, it pointed to the fact that the LCD screen was sub-optimal

for this purpose. This was not surprising, as the screen on Capshare was primarily intended for users to confirm what had been scanned in, and for flicking through the set of scanned images.

This finding emphasised to us that capture devices should better support the discussion of captured materials and designers should pay close attention to how information can be jointly viewed. In the case of digital cameras and Capshare their display screens – designed for capture confirmation – should be expanded and have wider viewing angles so as to better support information sharing, Moreover, capture devices could benefit from supporting the sending of information to projectors, monitors, printers and other display surfaces.

Capture to distribute

If the category "capture to discuss" puts the emphasis on information display technologies, "capture to distribute" points to the need for capture devices which are optimised to connect and communicate with other devices and infrastructures. For many capture devices, whether they be portable or desktop devices, users are required to go through several time-consuming steps in order to transmit captured information to other people.

For example, digital cameras are often designed with little support for easily sending the captured images to others. Part of the problem is uploading the contents of the camera to a networked environment such as a PC. Ease of doing this to some extent explains the success of the Sony Mavica camera on the market – this camera features a floppy disk drive built into the camera, facilitating copying photos from the camera to a PC. However, the taxonomy highlights that this does not go far enough. The problem is completing the *whole sending action*. This suggests that allowing send actions to be completed on the device itself would greatly improve usability. To investigate this further we have prototyped addressing emails and faxes directly on the Capshare. These are then sent when Capshare is put into a special "cradle" which allows quick connection and uploading of images to a networked PC. We have also prototyped sending emails and faxes directly from Capshare using a mobile phone, cutting out the need to use a PC entirely. Such solutions only arise from looking at the capture events within the context of their goals rather than focussing solely on the capture event in isolation.

Capture to archive

Capturing for the purpose of archiving is another interesting category because it has again very different implications for technology. Here, the capture opportunities we observed highlighted the need for easy connections to storage facilities, and for flexible search and retrieval mechanisms to operate on the captured information.

In the study there were many cases of subjects wanting to capture the output of meetings. This included flipcharts,

whiteboard scribblings, slides and notes. Often the reason for capturing them was rather unspecific, but rather it was done to add to the group archives just in case the information was ever needed. These examples and others showed that archived information can take many different forms, that it may need to be flexibly searched, and that the ability to add information to the material may be important in providing context at the time of retrieval.

The software provided with most digital cameras provide an example of how *not* to support archiving of materials. Folders of photos are often assigned arbitrary labels or numbers which provide no useful information when later searching or browsing. Keyword facilities are often cumbersome or time consuming to use, and there are no facilities to arrange by subject matter, or to flexibly organise and store the photos. Interesting ideas in this respect seem confined mainly to the research lab at this point [3].

Early versions of Capshare nominally met the requirements of archiving. As originally designed, Capshare placed its files into a standard Windows folder, but this did not allow documents to be easily browsed. By considering the kinds of archiving activities our subjects carried out, we suggested that the PC software needed to be expanded to allow documents to be kept ready to hand in one place, and to be easily browsed using thumbnails when searching. To support archiving, the software also needed to support documents being quickly moved into archive folders. At this point we were able to experiment with different kinds of off-the-shelf software to see how users reacted during search and retrieval tasks, and again, we could make a set of recommendations for improving archiving tasks.

Capture for task management
As a final example, we will consider capturing for the purpose of task management. This class refers to the capture of information for the purpose of organising one's tasks, and reminding oneself of future actions. Most of us have many ways in which we do this, such as the use of paper diaries, ToDo lists, and PDAs. In effect, such devices support "prospective memory" or memory for things we need to do in the future.

The diary data suggested that there are two important aspects to effective support of task management activities. First is the ability to allow users to capture information in an *ad hoc* manner. This is because people remember things they need to do at unexpected times and places. In fact, research shows that people are more likely to remember ToDo or action items when they are physically mobile [10].

Second is the ability to deliver appropriate reminders at the right time and in the right place. Our subjects showed that they recognised the importance of this. For example, Post-it notes and documents often were used as reminders, and these were placed in strategic places such as on computer screens or chairs so that they could not be missed. Most subjects carried reminding systems with them such as paper calendars or PDAs. Subjects also made a point of reviewing their lists and calendars on a regular basis so as not to miss important events.

The implications here are that support for task management is well suited to mobile technologies which allow ad hoc capture of information about ToDo items, and which can also deliver important reminders anyplace, anytime. The category of technology this points toward is PDAs, and these are obviously already successfully supporting capture for task management.

What was interesting, however, was the implications from this study for how PDA technology might be extended. One subject used the digital camera as a sort of visual ToDo list. She would regularly scan through her pictures on the back of the camera in order to remind herself of things she needed to do. For example, she took a picture of the photocopying room to remind herself to pick up some documents later. One reason she found this effective was that taking a picture was quicker than writing a note for herself. By using a digital camera as input, this quick action provided rich visual output as a reminder.

In addition to digital cameras, devices like Capshare also suggest new kinds of task management tools which could offer more than conventional PDAs. For example, one of our subjects photographed a letter he had received inviting him to a career evening at a local school. In taking the photo he said he wanted to be able to associate the document with his appointment in his electronic diary, so he would have the directions to the school handy when he was reminded of the event. This led us to think about ways in which scanned in documents could be usefully linked to appointment calendars.

Summary
The above discussion provides a demonstration of how the capture taxonomy we constructed can be used to impact the process of design. Simply by considering it we make a move to understanding the *goal* of capture. As we hope we have shown, many capture devices such as digital cameras have not been designed to optimise for goals; more typically they are optimised mainly for the capture task itself. The design of these devices may be improved by considering how well they current support the different actions in the taxonomy.

Considering the capture goals discussed in the taxonomy, it is striking how a set of activities in a single category tend to point in the same general direction, and toward particular classes of technology. Some categories point to emphasising better display technologies, some point to better communications, whilst others guide the designer toward storage and archiving, better software, or PDA-like devices.

Of course, this is not to say that capture devices of necessity must be specialised for one kind of goal to the exclusion of others. For example, some of the modifications that we have suggested for Capshare have shown that easier connection to a PC and well-designed software can provide an interface to allow users to easily email, archive, or print. In other words, some solutions better support a range of capture activities.

CONCLUSION

The primary aim of this paper was to unpack the goals behind capture activity so that we have a framework with which to understand how capture devices may be used, and thus to look for opportunities to invent new ones. To do so we used an innovative camera and diary methodology to naturalistically record both descriptive and quantitative data on capture activity in the context of people's working lives. The results show that people engage in a wide spectrum of different kinds of capture tasks. Further, for each different kind of capture "goal" that people may have, there are very different kinds of design implications. We illustrated this by showing how the framework can be used to inform the design of digital cameras as well as the Capshare document scanner.

Aside from these main findings, another point to make in closing is to highlight the success of this new diary technique. By using digital cameras to support diary-keeping, one can collect naturalistic data without the large overhead of observational approaches. Since the photographs are taken at the site of action, and interviews about the photos are carried out within a few days, subjects showed few recall problems when prompted by photos even on the relatively low resolution screen used on current digital cameras. This method also reduced the demand on the subjects themselves as taking pictures was easier for them than writing notes. This is not to say that this method is comparable to techniques such as ethnography, since as with all diary methods one is relying on the participation of the subjects. However, this technique is a useful addition to existing diary methods. We are enthusiastic about the further utility of this methodology as we move on to prototype and test some of our new ideas for mobile information appliances.

ACKNOWLEDGMENTS

BIBLIOGRAPHY
1. Adler, A., Gujar, A., Harrison, B., O'Hara, K., & Sellen, A.J. (1998). A diary study of work-related reading: Design implications for digital reading devices *Proceedings of CHI '98*, ACM, New York, 241-248.

2. Hindus, D. & Schmandt, C. (1992). Ubiquitous audio: Capturing spontaneous collaboration. *Proceedings of CSCW '92*, ACM, New York, 210-217.

3. Kuchinsky, A., Pering, C., Creech, M., Freeze, D., Serra, B., and Gwizda, J. (1999). FotoFile: A consumer multimedia organisation and retrieval system. *Proceedings of CHI '99*, ACM, New York, 496-503.

4. Lamming, M. & Newman W. (1992). Activity-based information retrieval. *Personal Computers and Intelligent Systems: Information Processing '92*, 68-81.

5. Minneman, S., Harrison, S., Janssen, B., Kurtenbach, G., Moran, T., Smith, I., & van Melle, W. (1995). A confederation of tools for capturing and accessing collaborative activity. *Proceedings of Multimedia '95*.

6. Newman, W. & Wellner, P. (1992). A desk supporting computer-based interaction with paper documents. *Proceedings of CHI '92*, ACM, New York, 587-592.

7. O'Hara, K., Smith, F., Newman, W., & Sellen, A. (1998). Student readers' use of library documents: Implications for digital library technologies. *Proceedings of CHI '98*, ACM, New York, 233-240.

8. Prosser, J. (ed.) (1998) *Image based research*. Falmer, London.

9. Sellen, A.J., & Harper, R.H.R. (1997). Paper as an analytic resource for the design of new technologies. *Proceedings of CHI '97* (Atlanta, GA.) New York: ACM Press, 319-326.

10. Sellen, A.J., Louie, G., Harris, J. E., & Wilkins, A.J. (1997). What brings intentions to mind? An in situ study of prospective memory. *Memory, Vol. 5, No. 4*, 483-507.

11. Whittaker, S., Hyland, P. & Wiley, M. (1994). Filochat: Handwritten notes provide access to recorded conversations. *Proceedings of CHI '94*, ACM, New York, 271-277.