MEMORY, 1997, 5 (4), 483–507

What Brings Intentions to Mind? An *In Situ* Study of Prospective Memory

A.J. Sellen
MRC Applied Psychology Unit, and Rank Xerox Research Centre, Cambridge, UK

G. Louie
Rank Xerox Research Centre, Cambridge, UK

J.E. Harris
CAMBIT The User Interface, Cambridge, UK

A.J. Wilkins
MRC Applied Psychology Unit, Cambridge, UK

In a naturalistic study, we aimed to uncover the relationship between thinking about and remembering intentions. Electronic badges allowed us to track the activities of subjects within their work environment. Over two weeks, subjects were asked to respond using a button on their badges (1) every two hours (Time task); (2) whenever they were in a particular room (Place task). In addition, whenever they thought about the task, they were asked to indicate this with their badges. Although subjects thought about the Time task more, they forgot to respond more often than in the Place task. In the Time task, there was a marked absence of thoughts about the task following successful remembering. When subjects remembered the Place task, thoughts increased with proximity to the target location. In both tasks, thoughts about intentions occurred more in places such as stairwells than in locations where people tended to settle. On the basis of these findings, possible mechanisms for prospective memory are discussed.

INTRODUCTION

Two important elements distinguish prospective memory from the kind of memory that psychologists usually study. First, prospective memory is largely

Requests for reprints should be sent to Abigail Sellen, Rank Xerox Research Centre, 61 Regent Street, Cambridge CB2 1AB, UK. Email: sellen@cambridge.rxrc.xerox.com

This project was funded by Rank Xerox Cambridge EuroPARC. We thank Mik Lamming, Mike Flynn, and Paul Dourish of EuroPARC for help with technical aspects of the experiment, and the employees of EuroPARC for their help and good nature during the running of the experiment.

self-cued in that one must ''remember to remember'' at an appropriate point in time. The second is that what is recalled is an intention or plan. This can be contrasted with the vast majority of psychological literature on memory which deals with experimenter-cued memory for events learned or experienced in the past.

Prospective memory, as has often been pointed out, remains largely unexplored empirically. In a fairly recent review (Kvavilashvili, 1992), only 24 experimental studies were cited. Yet in diary studies of everyday memory problems, failures to remember intentions constitute a substantial portion, if not the majority, of problems that people report (Eldrige, Sellen, & Bekerian, 1992; Terry, 1988). Thus, the practical significance of this kind of memory is certainly not reflected by our theoretical understanding of the topic. So far, the literature has had little to say either about the cognitive processes leading up to the remembering or forgetting of intended actions, or about the mechanisms that might be responsible for bringing intentions to mind.

One way of classifying tasks that involve memory for intentions is to distinguish *time-based* from *event-based* remembering tasks (a distinction also made by Einstein & McDaniel, 1990). Time-based tasks involve some kind of ''appointment keeping'' in that people are asked to do something at a specific instance in time, or to do something within some more general time window. Many everyday tasks are of this kind, such as remembering to take a pill at an appointed time, or remembering to telephone someone on their birthday. However, for many other everyday prospective memory tasks it is not necessarily the case that something needs to be done at a particular time. Rather, it is often the case that something must be done at a certain point during an activity, or in a particular place, or on meeting a particular person. Remembering to buy wine on the way home from work, or to post a letter on the way to get coffee, or to pass on a phone message when you next see someone are examples. These then, are dependent on the occurrence of certain events. Hence they can be called event-based prospective memory tasks. Of course, many prospective memory tasks such as catching a train are components of more complex plans that involve both time and event-based memory.

It may be the case that these two different classes of prospective task involve different processing mechanisms. Einstein and McDaniel (1990) have suggested that event-based tasks require less in the way of self-initiated retrieval processes than time-based tasks because they are externally, or, to use Craik's (1986) terminology, ''environmentally'' supported. Indeed, it is easier to characterise event-based prospective remembering in the same terms that we might use to describe many retrospective memory tasks—i.e. as a cue (such as an activity, person, or place) triggering an associated memory. For time-based tasks, the process by which intentions are brought to mind is less obvious. How does one make oneself aware of time, or remind oneself of time-based intentions, in the absence of obvious external cues? What is the nature of these self-initiated processes? To date, they remain a mystery.

The Relation Between Thinking About and Acting on Intentions

In an effort to learn more about the processes by which intentions are brought to mind, it may be instructive to examine the relationship between thinking about and carrying out intentions. One difference between thoughts about intentions and execution of them is whether or not the appropriate conditions for acting on the intention are met when the intention is "in mind". Knowing about the frequency with which intentions come to mind, in addition to knowing whether or not they are eventually carried out, may give important clues to the underlying processes.

Investigations in this area have been limited to time-based remembering tasks. In one study by Harris and Wilkins (1982), subjects were required to perform an action at pre-specified intervals, and the number of times subjects monitored a clock was measured. This study showed that the frequency with which subjects monitored the clock increased as the target time approached. In addition, it showed that lower monitoring rates were associated with failures to respond on time. In a different paradigm, Kvavilashvili (1987) measured whether or not subjects reported that they thought about an intention in the interval between asking subjects to hang up a phone and the moment at which they were supposed to remember to carry out that action. She found that the group who reported thinking about the intention were more likely to remember the target task. Both studies suggest that, at least for time-based remembering, it is important to think about the task in the interval between forming and acting on the intention.

One limitation of both studies, however, is that they required subjects to remember over very short intervals of time (from three to nine minutes). Tasks of this sort are likely to be quite different in nature to many of the everyday time-based tasks with which we are often faced. Remembering appointments or remembering to take medication tends to take place over a much longer time period of hours, days, or even weeks. In these short-term laboratory tasks, subjects may often have managed to maintain some level of awareness of the intention. At the very least, being in an experimental situation must have helped subjects to keep the intention in mind.

This also raises the unexplored issue of whether thinking about an intention in the interval between forming and acting on an intention is also important for successfully remembering event-based remembering tasks. If people are able to rely on the potency of an event as a cue for triggering recall of an intention, bringing the intention to mind may not be an important prerequisite for successful later performance.

How Do Intentions Come to Mind?

If thoughts about intentions are important for prospective remembering, the critical question is then "How do these thoughts come to mind?". Do we need to rely on external cues, or can we create some kind of "mental note" to ourselves to help us remember? Contextual cueing is one mechanism by which thoughts about intentions might be triggered. But in the absence of contextual cueing, what might these internal mechanisms be like? It seems reasonable to rule out mental rehearsal or simply keeping intentions "in mind" for anything other than short-term time intervals. Another possibility is that we create some kind of internal marker which our random internal "wanderings" come across from time to time. If this is the case, we might expect to see randomness in the occurrence of thoughts about intentions across time. Alternatively, it may be the case that people are capable of somehow internally controlling the degree to which they initiate thoughts about intentions. Thus, thoughts about intentions may be subject to some kind of self-prompting mechanism which can be used to regulate the frequency with which we think about intentions.

In a study of intentions recorded in diaries, Ellis and Nimmo-Smith (1993) provide evidence to suggest that people are less likely to recall intentions during activities that require attention. However, simply knowing that attention is a relevant factor does not help to pinpoint the cognitive processes responsible for intentions "surfacing" to consciousness. Attending to another activity may mean that we are less vulnerable to the influence of contextual cues, or it could interfere with internal sources of control triggering thoughts about intentions.

Experimental Approach

The experiment we will describe was exploratory in nature, having three main purposes: (1) to compare performance in a time-based and an event-based task; (2) to examine the processes leading up to and following performance on these tasks (specifically, to examine the frequency with which intentions are brought to mind); and (3) to explore the factors that might be responsible for bringing intentions to mind.

Towards these aims, the potential effects of both a rich array of external cues and of other competing activities in bringing thoughts about intention to mind emphasised the importance of carrying out this experiment in the context of everyday life, and over longer periods of time than a laboratory study would permit. However, one of the major obstacles up until now has been the difficulty of tracking and recording memory performance in situations outside the laboratory.

Wilkins and Baddeley (1978 were able to study prospective memory *in situ* to some extent by asking subjects to push a button at appointed times on "button boxes" over the course of seven days. These boxes were portable electronic devices which recorded the time of day whenever the button was pressed.

Unfortunately, the button boxes provided no information (other than the time) about the context in which intentions were remembered. In this study, we were able to collect contextual data by using a more sophisticated portable technology called "Active Badges"[1] (Want, Hopper, Falcao, & Gibbons, 1993). Badges provide a record of where badge-wearers are at all times during their working day. These contextual data are collected unobtrusively and automatically, and the collection can occur over long periods of time.

This technology allowed us to design a study in which we could systematically investigate prospective memory performance outside the laboratory, in both a time-based and event-based task. We were interested in tasks to be carried out without the deliberate use of external memory aids, so that people were solely dependent on their own mental resources.

METHOD

Technology

The Active Badge system is installed in Rank Xerox's Research Centre in Cambridge which houses approximately 25 employees. Badges are lightweight devices which clip to one's clothing (Fig. 1) and emit an infra-red beam approximately every 20 seconds. These infra-red signals are picked up by special sensors installed throughout the building, there being at least one sensor in every room, corridor, and stairwell (with the exception of the lavatories). The sensors are constantly polled via a computer network, and information about which badges are registered at each sensor is recorded in a central database. Specialised computer software then generates an electronic "diary" for each badge-wearer by extracting the relevant data from the central database about that badge-wearer's activities over the course of the day. Thus, as badge-wearers move about the workplace, the places that they enter, the time they spend in each place, and the presence of other badge-wearers is automatically recorded. Badges also possess a small button. If the button is pressed, the moment and context in which the button is pressed is also recorded by the badge sensing system (see Lamming & Newman, 1992, for more details).

Subjects

Subjects were recruited from within the building, and only those who planned to be in the building for the majority of the four weeks during which the experiment was being conducted were eligible. Only 14 people met these criteria, all of whom agreed to participate. Five of these people had

[1]Active Badges were originally developed at Olivetti Laboratories in Cambridge. They have been used at Rank Xerox EuroPARC as part of a larger effort to build applications to support human memory (see Lamming et al., 1994, for more details).

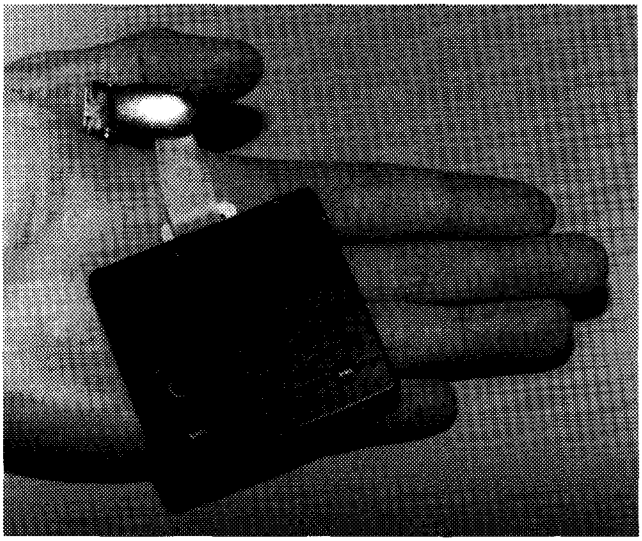

FIG. 1. An ''Active Badge'': a wearable device that emits an infra-red beam enabling the automatic recording of people's activities.

administrative jobs, while the other eleven were involved in research, either as senior research scientists (five), students (three), or consultants (two). Eight subjects were male and six female.

Tasks

Subjects were asked to carry out two kinds of memory tasks, each task for a period of five consecutive working days. It was stressed that for both tasks, subjects were not to deliberately use any kind of external memory aid (e.g. notes to themselves, electronic reminders), nor even to discuss the tasks with anyone else in the building (as others might deliberately or involuntarily remind them of the tasks).

Both tasks involved two simple activities: (1) remembering to press the button on their badge either at a pre-specified time (Time task) or at pre-specified place (Place task); and, (2) pressing the badge button whenever they thought about the task, regardless of where they were or what time it was. Button-pressing consisted of aiming the badge at the nearest wall sensor, and pressing the button three times in quick succession, henceforth called ''triple-clicking''. Button pressing at the target time or place consisted of triple-clicking *twice* leaving about two seconds in between the two sets of triple-clicks. This was to distinguish clicking at the target time or place from simply thinking about the intention. Subjects were instructed to triple-click only *once* whenever they thought of the task. Reporting thoughts in this way is of course itself a prospective memory task, but one that we thought unlikely to introduce systematic error into the primary task. In both tasks, subjects were asked to try to click their badges discreetly so as not to remind other subjects who might be nearby of the experimental tasks.

Time Task. During the Time week, subjects were asked to triple-click their badge button twice at prescribed ''target times'' throughout the day. Each subject was randomly assigned to one of four different target time schedules, and followed the same schedule every day. The first target time of the day occurred at one of four times, either at 7:30, 8:00, 8:30, or 9:00 a.m. Subsequent target times then occurred at successive two-hour intervals. (Different schedules were assigned to different subjects to help prevent people who might be nearby from cueing each other.)

Subjects were asked to triple-click twice somewhere within a five-minute interval from each target time to five minutes past that time, *given that they were in the building.* If they had not arrived at work for a particular target time, or if they were not in the building when a target time interval occurred, they were informed that this would not be counted as an opportunity to respond.

Thinking about the Time task was indicated by asking subjects to triple-click once whenever they thought of the task, regardless of what time it happened to be when this occurred.

Place Task. During the Place week, subjects were asked to triple-check twice whenever they found themselves in an area known as the ''Commons''. The Commons is a public area containing seminar facilities, kitchen facilities, mailslots, and the library. It is therefore a room that employees visit frequently and for a variety of reasons. Each separate visit to the Commons was counted as an opportunity to respond, and subjects were told that triple-clicking twice at any time during each visit would count as a successful response.

As with the Time task, during this week subjects were also asked to triple-click once to indicate that they had thought of the Place task, regardless of where in the building they happened to be when this occurred.

Design and Procedure

The experiment took place over a total of four weeks. During the first and third weeks, subjects were simply instructed to wear their badge every day. During the second and fourth weeks, subjects carried out the memory tasks. Half of the subjects carried out the Time task in the second week and the Place task in the fourth week, and the other half did the reverse. Thus task was a within-subject variable.

No tasks were assigned in the first week simply because some of the subjects did not ordinarily wear badges, so this week provided an opportunity for these people to get used to wearing them, and for people who were not subjects to get used to seeing them.[2] No task was assigned in the third week in order to minimise possible carry-over effects between the two tasks.

During the first and third weeks, all subjects were automatically sent a daily reminder to wear their badges via electronic mail. During the Time and Place weeks, subjects were mailed either the Place or Time instructions (whichever were relevant) on Monday morning only. On reading these instructions, they were asked to electronically mail the experimenters to indicate exactly what time they had read them, and then to immediately delete them so that they would not serve as a further reminder.

At the end of the four-week period, all subjects were given a structured interview designed to uncover any possible technical problems with the experiment, but more importantly to elicit subjects' comments about the remembering tasks themselves.

RESULTS

Of the original fourteen subjects, two (both with administrative jobs) had to be discounted from the analysis. One subject was ill for the majority of the four weeks. Another decided to opt out of the experiment due to personal circumstances unrelated to the experiment. Fortunately, both subjects were from different groups, so the counterbalancing of task order was not disrupted. For the remaining 12 subjects, activity diaries for each day of the experiment were generated.

Overall Performance Measures

Henceforth, remembering a triple-click twice while in the target place (in the Place task) or during the target time (in the Time task) will be referred to as a "hit", and forgetting to do so will be called a "miss". Each target time interval that occurred while subjects were recorded as being in the building (in the Time task), and each recorded occurrence of visiting the target place (in the Place task) will be called a "target opportunity".

Table 1 shows the overall statistics for the Time and Place tasks. "Percent correct" refers to hits as a percentage of target opportunities encountered over the course of the week.[3] "Number of thoughts" refers to the number of times subjects indicated that they thought about the task (by one triple-click) when

[2]That being said, wearing a badge is not an unusual sight at EuroPARC as many people wear them for purposes other than experiments such as this.

[3]The interviews revealed that some of the clocks and watches subjects used were likely to be out by as much as three minutes, resulting in the fact that even the most conscientious subjects were sometimes consistently early or late. Therefore, in the case of the Time task, if subjects were three minutes early or late, and triple-clicked twice, this was still counted as a hit.

TABLE 1
Time and Place Tasks

	Type of Task			
	Time		*Place*	
% Correct	33.0	(29.8)	51.9	(34.1)
Number of Thoughts per Day	4.4	(3.7)	2.7	(2.8)
Number of Target Opportunities per Day	3.7	(1.2)	5.7	(3.9)
Number of Thoughts per Target Opportunity	1.4	(1.1)	0.7	(0.8)

Percent correct, mean number of thoughts, and mean number of target opportunities, for Time and Place tasks (standard deviations in brackets).

they were *not* in the target place (during the Place week) or when it did *not* occur during the target time interval (during the Time week).

The first point of interest is the much lower success rate in this time-based task than that found in Harris and Wilkins' (1982) time-based task. Harris and Wilkins' subjects responded within the target time interval 86.5% of the time at three-minute intervals, and 88.8% of the time at nine-minute intervals. That only 33% of responses were remembered in this time-based task is perhaps understandable, given that it was performed in a much richer environment, and over much longer intervals.

There was a significantly higher percentage of hits in the Place task than in the Time task [$F(1,10)=7.09$, $P<.024$]. However, subjects indicated they thought about the Place task less often than they thought about the Time task when measured on a per day basis [$F(1,10)=7.64$, $P<.020$]. Given that there were more target opportunities for subjects to perform the Place task than the Time task on a daily basis [$F(1,10)=8.16$, $P<.017$], this difference in number of thoughts is accentuated when calculated per target opportunity [$F(1,10)=9.91$, $P<.010$], being on average twice as frequent in the Time task as in the Place task.

The overall performance data indicate that the event-based task was an easier task to remember, and required fewer thoughts than the time-based task. There was large inter-subject variability: performance in the Time task ranged from 0 to 65%, and in the Place task from 20% to 92%. Nonetheless, with one notable exception (see later), performance for each subject was consistently higher in the Place task than in the Time task. When interviewed, these 11 subjects confirmed that the Place task was easier. Most gave similar reasons for this, claiming that in the Place task, they could rely on the activity of going to or entering the Commons as something that would trigger the remembering of the intention to badge click. The Time task, they said, was more difficult as there was no such reliable cue available. The following comment from one of these subjects is typical:

> The Place task was easier. Looking at the box [badge sensor] when entering the Commons reminded me. Often, just going to the Commons was enough to remind me. The Time task was done by keeping it in my brain. I would concentrate on the one coming up. It [the Time task] was more of a worry because it was harder and I had to make myself be aware of it.

The notable exception was an administrative assistant for whom percent correct was higher in the Time task than in the Place task (41.7% vs. 13.7%, respectively). When interviewed, she confirmed that, for her, the Time task was easier. Interestingly, she was also the only subject to claim that she had a very routine day. She said that she used the close correspondence between time of day and daily activity as a strategy for remembering the Time task. For example, she knew that her 9:30 target opportunity occurred when she regularly sorted the post, and her 3:30 target time occurred when she laid out the tea things. By associating her target times with predictable activities, she essentially transformed the time-based task into an event-based task, providing further evidence that reliance on some kind of contextual cue facilitated remembering intentions.

In general we can consider that the people in this particular workplace have less fixed schedules, and are less constrained by time, then in other kinds of workplaces. It is part of the ethos of the laboratory that people ''make their own schedules'' and monitor their own work, starting and ending their working day as they see fit. Although this style of working may be different from other kinds of organisations, it is somewhat typical of other organisations that employ a high concentration of ''knowledge workers'' (or professionals who are employed to understand and make sense of a body of knowledge, e.g. Kidd, 1994). Thus we may assume that such people have less structured days in general. Exceptions include administrative staff who are more likely to keep ''typical office hours''. In the case of the exceptional subject mentioned earlier, it is probably no coincidence that the fact that her day was the most structured also reflected the fact that she was the most junior administrative person, and therefore the one responsible for the activities that did require her to carry out certain activities at specific times.

Effects of Other People

One possible reason for the Place task being easier is that the Commons is, by nature, a place where people meet, and thus the higher likelihood of seeing others wearing badges in the target place might trigger remembering. Indeed, one of our concerns was that people would see each other badge-clicking, despite our instructions to subjects to click their badges discreetly. To check this, we counted the number of instances of ''group clicking'' as indicated by people in the same room clicking within seconds of each other (clicking being defined as either hits or thoughts). In all, there were very few occasions on which this happened, occurring in 6 group situations and accounting for 19 responses in the Place task, and occurring in 4 group situations and accounting for 11 responses in the Time task.

Because everyone in the building wore a badge during the course of the experiment, whether or not they were actually participating as subjects, we were able to check further whether the number of people present influenced the likelihood of a subject thinking about or remembering the target task. To do this we calculated the expected number of clicks alone versus in the presence of other people, on the assumption of constant probability, by multiplying the proportion of time spent alone by the total number of hits plus thoughts for each subject. We then compared this to the actual number of clicks alone versus with others for each subject. We found no significant difference between the observed numbers and what one would expect under a uniform distribution for the Time task [$F(1,11)=2.77$, $P<.125$] or for the Place task [$F(1,11)=2.40$, $P<.150$]. On the basis of this analysis, we can rule out any significant effects due to the presence of other people.

Effects Over Time

Performance on the memory tasks in terms of percent correct per day did not change significantly for either the Time or Place task over the course of the week. Nor were there any significant Order by Task interactions for percent correct (i.e. whether the Time task or Place task was performed first). However, the average number of thoughts per day did vary over the course of the week [$F(4,40)=4.48$, $P<.004$]. Post hoc tests of the differences among means using the Peritz procedure (.05 level, Toothaker, 1991) showed a decrease in the daily frequency of thoughts over the course of the week in both tasks. In the Time task, the frequencies on Monday and Tuesday were significantly higher than Thursday's mean frequency (the frequency increased slightly but not significantly on Friday). In the Place task, the frequency on Monday was significantly higher than the means for Wednesday, Thursday, and Friday.

One interpretation of these data is that subjects were becoming more efficient in their remembering, in the sense that they maintained a consistent level of performance on the memory tasks while having to think about the task less over the course of the week. In other words, they may have learned that thinking about the task so much was not necessary as they become experienced with it. Another possible interpretation is that subjects were simply becoming less motivated to indicate their thoughts as the week wore on.

There was also a significant Order by Task interaction for frequency of thoughts, whether measured as thoughts per day, or as thoughts per target opportunity [$F(1,10)=8.33$, $P<.016$, and $F(1,10)=6.29$, $P<.031$, respectively]. This showed that the number of thoughts in the Time task was higher if the Time task was performed first, as opposed to performed second. Similarly, the number of thoughts in the Place task was higher if the Place task was performed first, rather than second.

These analyses suggest that whatever the reason for the frequency of thoughts decreasing over time (whether motivational or otherwise), these factors operated not only over the course of the week, but also over the course of the experiment. Fortunately, there was no indication, statistical or otherwise, that these factors operated differently for the two tasks, so interpretation of the data is not confounded by these results.

Distribution of Thoughts in Relation to Target Opportunities

An analysis was carried out of how thoughts were distributed in the periods leading up to and following the target opportunity in both tasks.

Time Task. The frequency of thoughts for all subjects in relation to the beginning and end of the target time interval is shown in Fig. 2, summed over two-minute bins. For each target time interval, thoughts occurring in the hour preceding and the hour following the interval were plotted. Each two-hour interval between target times was divided in half, and thoughts were plotted in relation to the preceding or succeeding target opportunity, depending on whether they occurred in the first or second hour of that interval. They were also classified according to whether they preceded or succeeded a hit or a miss.

An analysis of thoughts based solely on frequency measures is potentially misleading, however, due to the fact that there is variation in the total number of opportunities during which thoughts can be registered by subjects. As is shown by the solid lines in Fig. 2, the number of opportunities during which thoughts can be expressed varies across time. This is because the total time subjects spent in the building varied across time. In addition, there were fewer opportunities overall for hits versus misses. This is because there were fewer hits than misses in the Time task overall, and thus fewer opportunities on which thoughts associated with hits could be expressed. Thus, the relative frequency of thoughts is not necessarily the best estimate of the true likelihood of thoughts at any moment in time. For the purpose of statistical analysis, an estimate of the probability of thoughts in relation to the target time interval was calculated by dividing the frequency of thoughts per two-minute bin by the average number of opportunities available per bin.

Analysis of variance was computed on measures of thoughts/opportunity for the within-subject variables of Response Type (hit/miss) and Relation to Target (before/after). There were no significant main effects, but there was a significant Response Type by Relation to Target interaction [$F(1,11) = 12.22$, $P < .005$]. Post hoc means tests (Peritz, .05) showed that the probability of thoughts after hits was significantly lower than the probability of thoughts before hits, and was

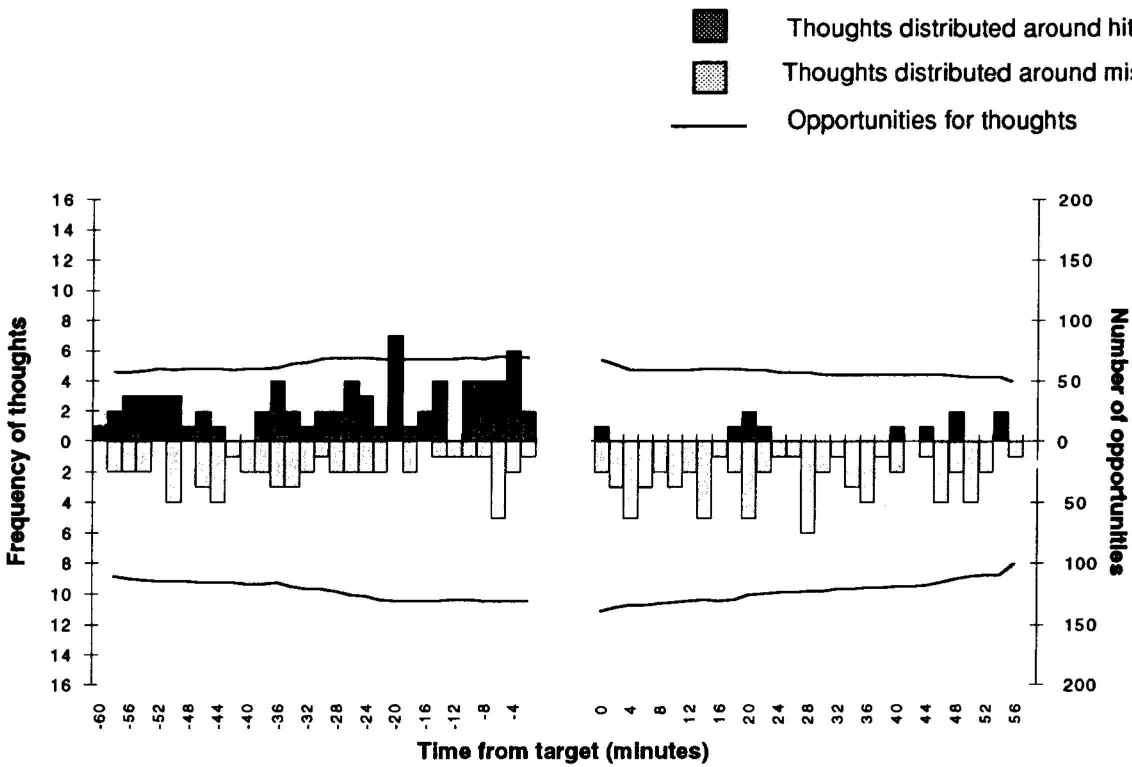

FIG. 2. Time task. The distribution of the frequency of thoughts leading up to and following the target time interval for all 12 subjects (scale on left axis). Thoughts shown on the upper half of the figure correspond to thoughts distributed around hits, while those on the lower half correspond to thoughts distributed around misses. The solid lines indicate the number of opportunities on which thoughts could be expressed (right axis).

significantly lower than the probability of thoughts before and after misses. No other differences were found.

Place Task. The frequency of thoughts for all subjects in relation to the target place is shown in Fig. 3. Thoughts are plotted up to the time at which subjects entered the Commons, and from the time at which they left the Commons. As in the Time task, the interval between target opportunities was divided in half: thoughts in the last half of this interval were plotted as leading up to the succeeding target opportunity; thoughts in the first half of the interval were plotted as thoughts following the preceding target opportunity. Unlike the Time task, however, the interval between target opportunities was variable. Thus, the number of opportunities on which subjects could express a thought varied with the size of this interval as well as with the amount of time subjects spent in the building during the day. The solid lines in Fig. 3 show how these opportunities vary across time. The distribution of opportunities also varies according to whether the responses were hits or misses, but as about half of the targets were hit, this factor does not have a large effect on number of opportunities.

As in the Time task, the probability of thoughts was estimated by dividing the frequency of thoughts by the number of opportunities available. However, because the number of opportunities is substantially lower at the extreme ends of the time scale, the estimates are likely to get increasingly less accurate towards those extremes.

Analysis of variance was again computed on the probability measures for the within-subject variables of Response Type (hit/miss) and Relation to Target (before/after). As in the Time task, there were no significant main effects, but a significant Response Type by Relation to Target interaction [$F(1,11)=10.36$, $P<.008$] was found. Peritz tests showed that the probability of thoughts before hits was significantly higher than the probability of thoughts after hits, and was significantly higher than the probability of thoughts before and after misses. No other differences were found.

Figure 3 shows that this interaction is due to the large increase in thoughts occurring just prior to entering the target place. Indeed, a more detailed examination of the data shows that 71% of the responses in the two-minute interval just prior to entering the Commons occur within 30 seconds before entrance. They therefore occurred mostly in the stairwell or corridor on the way to the Commons. Proximity to the target thus appears to trigger thoughts, which is not the case in the Time task.

Interestingly, Fig. 3 shows that there is also a peak in thoughts (albeit smaller) on exiting the Commons, but only following a missed response. These thoughts all occur within five seconds of leaving the target place, and almost all of them occur in the stairwell. Physical transition into or out of the target place appears to have increased the likelihood of thinking about intentions in the Place task.

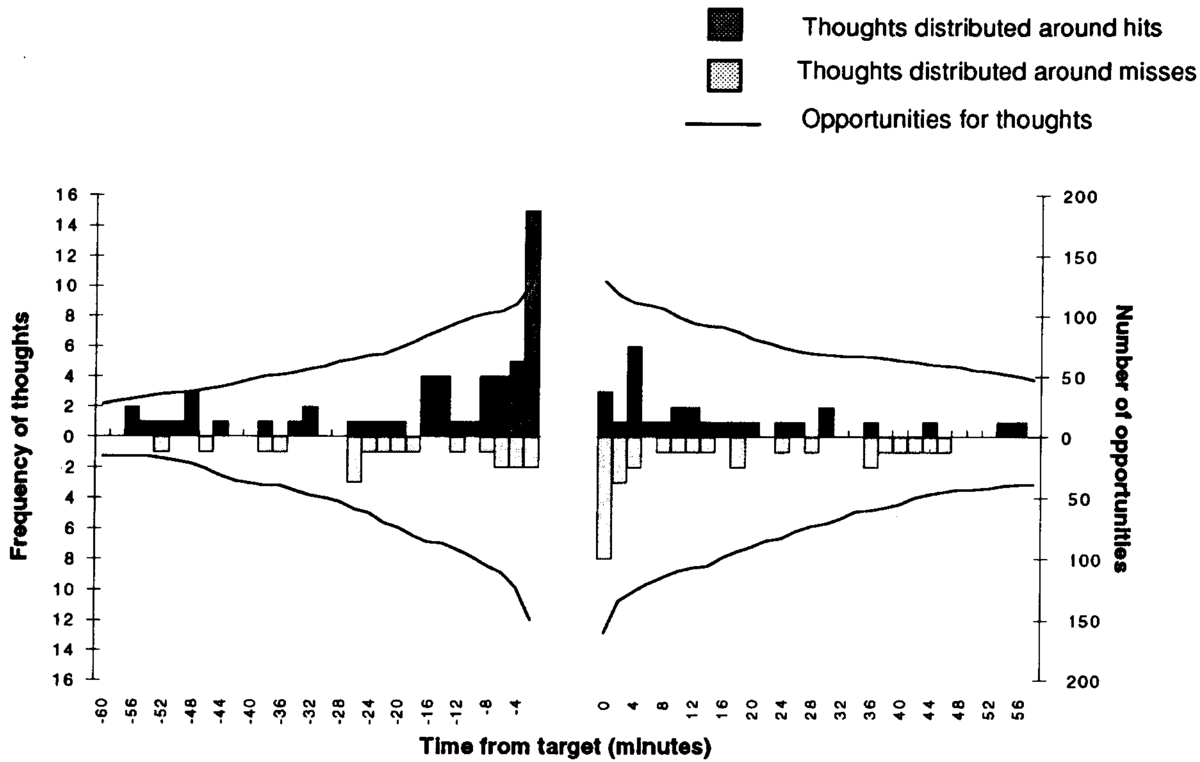

FIG. 3. Place task. The distribution of the frequency of thoughts leading up to the entrance and following the exit of the target place, for all 12 subjects (scale on left axis). Thoughts shown on the top half of the figure correspond to thoughts distributed around hits, while those on the bottom correspond to thoughts distributed around misses. The solid lines indicate the number of opportunities on which thoughts could be expressed (right axis).

The Place task shows quite a different pattern of thoughts over time than the Time task, therefore. Not only is there an effect of proximity to the target place, there is no marked absence of thoughts immediately following a hit, as was found in the Time task. Rather, with the exception of thoughts in immediate proximity to the target place, the frequency of thoughts appears to be more evenly distributed in time whether in relation to a hit or a missed response.

In Figs 2 and 3 it is also of interest that the bars representing thoughts just before a miss are not empty. This indicates that, in both tasks, people sometimes had thoughts right up to the last bin before the target opportunity and yet still failed to respond during that opportunity. This is one result that is consistent with the findings of Harris and Wilkins (1982).

Distribution of Hits Within Target Opportunities

Not only was approaching the target place associated with an increase in thoughts about intentions in the Place task, but the act of entering the Commons also appeared to be associated with triggering hit responses within the target place. Figure 4 shows both the Time and the Place tasks plotted as a percentage of time spent in the target interval or place. Percentage rather than absolute time is shown because the duration of visits to the target place in the Place task was highly variable (visits ranged from .08 to 230 minutes, the mean being 14 minutes per visit).

Figure 4a shows that, for the time task, hits are fairly evenly distributed across the target time interval. However, Fig. 4b shows that in the Place task, the frequency of hits peaks on first entering the Commons, declining steeply thereafter. This was the case even thought subjects were told that they could make their response at any time during each visit to the target place. In absolute terms, 68% of hits were made within 10 seconds of entering the Commons. No peak was observed immediately before exiting the Commons, although, as previously observed, there was an increase in thoughts immediately *after* exiting.

The Relation Between Location and Thoughts

So far the results have provided evidence that making a transition in or out of the target place is associated with an increase in thoughts about the Place task. This raises the question of whether being in a state of transition among other locations in the building also gave rise to an increase in thoughts about intentions.

In order to assess the influence of being in transition on the frequency of thoughts, all of the locations (rooms, corridors, stairwells, etc.) in the building were rank ordered according to the average length of a visit to each location over the two weeks in which the memory tasks were performed. Locations that, on average, were only visited or passed through for very short periods of time

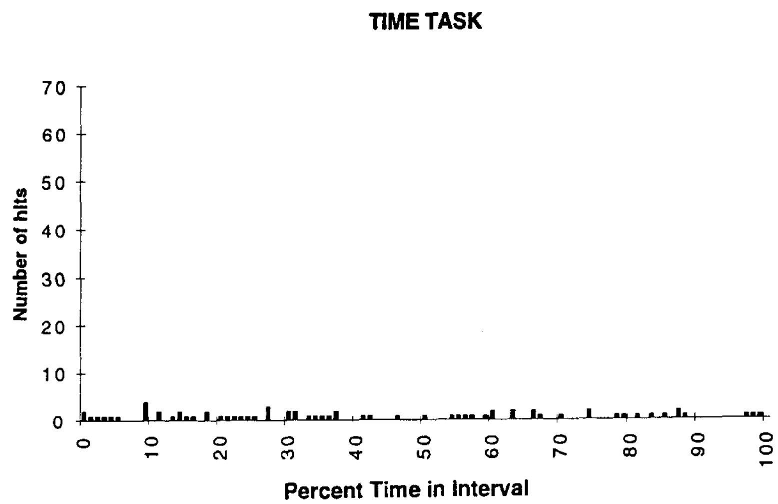

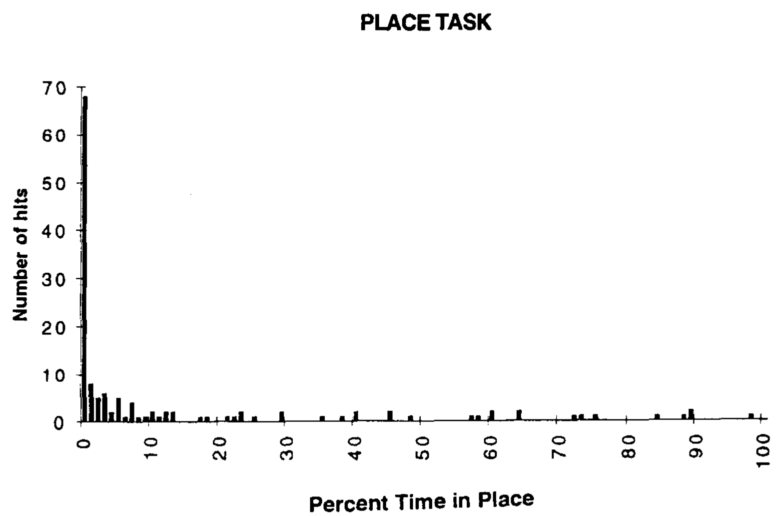

FIG. 4. The distribution of hits within (a) the target time interval (Time task) and (b) the target place (Place task). Frequency of hits are shown as a function of percentage of total time spent in target time interval or target place.

(''transitional locations'') can be contrasted with those that were visited for long periods of time (''settling locations''). Locations visited for an intermediate duration fall in between these extremes and can be termed ''intermediate locations''.

Figure 5 shows the locations within the building ranked in increasing order of average duration per visit to that location. Some locations, such as ''other people's offices'', are the result of collapsing several locations into one category. As we might expect, locations such as corridors and stairwells fall at the ''transitional'' end of the scale, whereas locations such as one's own office and the conference room fall at the other end, being places where people are more likely to settle for long periods of time. Intermediate locations include rooms in the building called ''pods'' which are common areas shared by groups of offices containing equipment such as photocopiers and printers. Also shown in Fig. 5 is the average number of thoughts per hour spent in each location as a function of location, for both tasks.

Figure 5 shows that the highest rates at which thoughts occur in both tasks cluster around the transitional locations (i.e. in the stairwell, lobby, and corridors). That an unusually high number of thoughts occurred in transitional locations as opposed to intermediate or settling locations is confirmed by statistical analysis. A binomial test was used to compare the number of thoughts in transitional places with the number of thoughts that occurred in non-transitional places (intermediate plus settling places). The observed number of thoughts was compared to the number expected on the basis of a uniform probability of thoughts across time. The resulting P values were then combined using Edgington's (1972) method of combining probability values from independent experiments, treating each subject as an independent experiment. Subjects with a total of five or fewer thoughts overall were excluded from this analysis (two subjects in the Time task, and four in the Place task).

The hypothesis of equal probability of thoughts in transitional versus non-transitional locations was rejected for the Time task, whether subjects did the Time task first [$P < .0001$] or second [$P < .023$]. It was also rejected for the Place task [$P < .0001$] whether they did the Place task in the first or second week.

One could argue that we might expect this result in the Place task, given that we already know that proximity to the Commons tended to lead to an increase in thoughts. Locations proximal to the Commons tended to be places like stairwells and corridors. However, if we remove thoughts occurring within 30 seconds of entering or leaving the Commons from this analysis, the results still show that the highest rates cluster around the transitional locations. The combined result of binomial tests on the data with these thoughts removed is also still significant for the Place task [$P < .0001$] whether it was performed first or second. It is also worth noting that subjects who performed the Time task in the first week also indicated increased thoughts in transitional locations. Therefore, this effect in the Time task cannot be attributed to any carryover effects from having

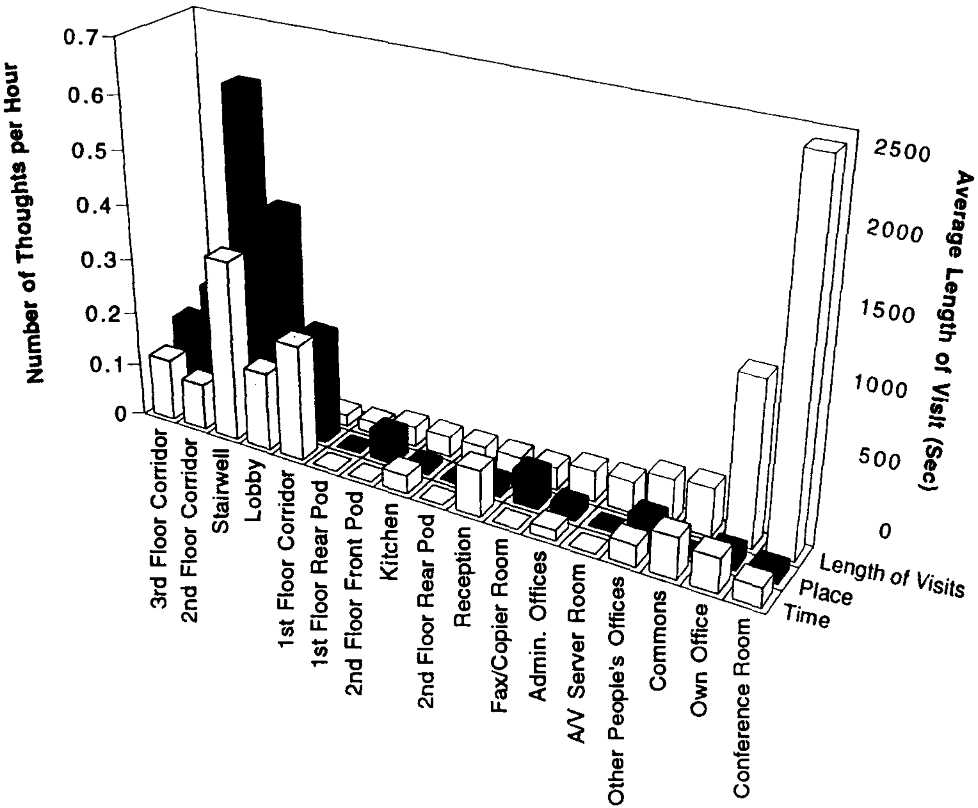

FIG. 5. Number of thoughts per hour as a function of location in the building for both the Time task (grey bars) and Place task (dark bars). Locations are rank ordered by the average length of a visit to each location by subjects during the experiment (shown by white bars).

performed the Place task first. There is therefore no reason to assume that the Commons took on any special significance for this group of subjects, eliciting more thoughts on the way to and from that room, perhaps accounting for increased thoughts in stairwells and corridors. Thus, in both tasks, this finding is not simply an artifact of proximity to the Commons.

DISCUSSION

This paper began by drawing a logical distinction between prospective remembering tasks that must be carried out on the basis of time, and those that must be carried out on the basis of the occurrence of some event. Different patterns of results for time-based and event-based tasks found in this experiment support this distinction. In particular, the fact that the place-based task was easier to carry out than the time-based task suggests that association with a contextual clue (in this case the activity of going to a specific place) facilitates the remembering of intentions. Not only was performance in the Place task better than the Time task, but both the objective data and reports from the interviews indicate that there was less need to think about the Place task in the interim, as compared to the Time task. One interpretation of this is that reliance on a contextual cue decreases the need to periodically bring the intention to mind, or to "worry" about the to-be-remembered task. Converging evidence for the facilitory effects of contextual cueing is given by the one subject who performed set activities at the same time every day. Time could therefore be cued by activity, and performance on the Time task was, for her, easy compared to the other subjects who had unstructured daily routines.

The role of contextual cueing in the Place task is further illustrated by the fact that an increase in physical proximity to the target place corresponded to an increase in the likelihood of calling an intention to mind. Temporal proximity to the target time interval in the Time task had no such analogous effect in terms of an immediate increase before the target time. More specifically, in the Place task, both the frequency of thoughts and the frequency of hits while in the target place reveal which aspects of the target event were triggering recall of intention. It appeared that approaching and, in particular, the act of entering the target place was the salient part of the event provoking recall of the intention. Although making the transition into the target place was found to be an effective cue for remembering, no similar effect was found on exiting the target place following a successful response. However, if the intention to badge-click in the target place was forgotten, an immediate increase in thoughts was registered just after exiting the target place. It does appear, then, that for some subjects, making the transition out of the target place acted as an effective cue for realising that they had forgotton to respond.

A different kind of underlying mechanism for prospective memory was indicated by the pattern of occurrence of thoughts in the Time task. First, the higher frequency of thoughts in the Time task suggests that subjects were making more of an effort to periodically bring intentions to mind than in the Place task. That subjects were exerting some kind of internal control over the frequency of these task-related thoughts is also suggested by the interview data. When asked what caused them to remember or think about the Time task, most of them said they had to ''make'' themselves aware of the task or had to try to ''keep the task in mind''. No such attempts at a conscious effort to remember were reported by subjects for the Place task. Instead, subjects cited a range of external cues as being responsible for reminding them of their prospective task, not limited to the Commons. For example, seeing badges or hearing about badges were frequently mentioned as cues that caused reminding.

Second, in the Time task, subjects showed a marked decrease in the frequency with which thoughts came to mind following a successful response. Thus it appeared that after remembering the Time task, subjects were able to put aside task-related thoughts. Perhaps, following a missed response, they were unable to do so because of an increased amount of ''worry'' which carried over to the next target interval.

If thoughts came to mind randomly, or if they depended solely on being triggered by task-related external reminders, no systematic difference in the frequency of thoughts before and after remembering would be expected. If, on the other hand, time itself served as a salient and specific cue, then one would expect, as in the Place task, an increase in thoughts with proximity to that cue. Instead, the results suggest a mechanism that is sensitive both to performance on the memory task, and also to the passage of time in a rather coarse-grained way. In this sense, the results differ from the findings of Harris and Wilkins' (1982) laboratory study in which they found that the frequency with which subjects monitored a clock increased with proximity to the target time, mainly in the last minute before the target time. No such fine-grained proximity effect was observed in this experiment.

One possible reason for this difference is that this is due to differences in the granularity of the response measures in the two situations. Clock monitoring may be a more sensitive measure of the frequency of thoughts about intentions than badge-clicking. In other words, subjects in this experiment may have decided to register only one badge-click response for thoughts that occurred very close together in time. However, even if this were the case, we might still expect to see an overall cumulative increase in thoughts immediately preceding the target time.

An alternative explanation is that the time-based tasks in this study and in Harris and Wilkins' study were fundamentally different. It is not necessarily the case that subjects in the laboratory study were monitoring the clock every time they thought about the task. It could be that subjects more or less continuously held the task in mind, simply making the decision to monitor when their uncertainty about the time reached some critical level. Sensitivity to the target time may only occur under conditions where the intention is never truly "out of mind". In an experiment where a response is required over a long period of time and a real-world context, it appears that we must appeal to different mechanisms which are not so finely sensitive to an impending target time.

Two Classes of Mechanism

Before drawing some general inferences, we must add that caution is necessary in generalising the results on four counts:

1. The Time and Place tasks differed in the complexity of the task demands in so far as the Time task required the remembering of several *different* targets whereas only one target was involved in the Place task. This draws attention to three components of prospective memory tasks: (1) remembering *what* to do; (2) remembering the critical conditions under which it is to be done; and (3) recognising the retrieval cues as such when they occur. In terms of the present experiment, subjects were required to: (1) press the button three times; (2) to do so when in the Commons or at specific two-hour intervals; and (3) recognise the Commons or the time of day as the trigger. This implies that the Time task may in fact impose a greater load on retrospective memory in that the critical conditions to be remembered are more complex.

2. As discussed earlier, thoughts about intentions may occur more frequently than subjects are prepared to report them, and this under-reporting may be more marked when thoughts are frequent. Although under-reporting may explain our failure to show an increase in thoughts immediately prior to the target interval (cf Harris & Wilkins, 1982), the fact that the same reporting method was used in the Place task, which *did* find an increase immediately prior to the target, casts doubt on this explanation. All in all, the amount of under-reporting might be expected to increase with the number of thoughts to be reported. If this was the case, it only serves to strengthen the difference between the Time and Place tasks.

3. The time-based remembering and event-based remembering were not directly comparable in terms of factors such as the frequency with which remembering was required. However, it is the pattern of results within each task that is of interest, rather than the direct comparison of absolute numbers across tasks. These patterns are dramatically different in the Place and Time tasks. Precisely what it is about the two tasks that brings these differences about needs further analysis under different experimental conditions.

4. We used only one type of time-based remembering task and one type of event-based remembering. One could argue, for example, that the Place task is very different from other kinds of event-based tasks, such as remembering to give someone a message when they enter the room. Going to a specific place is under one's control, which is not necessarily the case in passing on a message to someone. Factors such as this may affect prospective memory performance.

With these cautions in mind, the different patterns of results for the two tasks suggest two fundamentally different classes of mechanism responsible for causing the remembering of intentions. The first kind of mechanism involves prompting from the ''outside–in'', in that intentions are triggered by external cues. This external prompting is different from an experimenter prompting a subject to remember—subjects in this experiment still had the task of ''remembering to remember'' when faced with a contextual cue. However, as McDaniel and Einstein (1993) have shown, measurable aspects of such cues, such as their familiarity or distinctiveness, can affect the extent to which they have the power to evoke remembering.

The second kind of mechanism postulated is presumed to be from the ''inside–out'' in that it is up to the individual to initiate thoughts about intentions, prompting them to then attend to external conditions (i.e. checking the clock). The results of this experiment suggest that this kind of self-prompting is under the control of the individual in more than one way—that the frequency of producing self-initiated thoughts can be controlled, to some extent, over time. Of course, many time-based intentions, such as catching a train, will be part of more complex-goal directed hierarchies, an issue discussed by Harris (1984). In such cases, self-prompting may be triggered by thoughts or actions associated with higher-level goals. The kind of time-based task in this experiment was much more akin to taking a pill, in that the action required is a separate goal in itself and not part of an ongoing series of activities. Presumably this is one reason why such simple time-based activities are so difficult to remember.

Having distinguished between these two kinds of mechanism, we are not claiming that the Time and Place tasks operated according to completely different principles. Indeed, when interviewed, subjects said that external events such as seeing or hearing mention of things related to time, such as clocks or diaries, acted as reminders in the Time task. Badge-related things were also cited as reminders for both tasks. On the basis of our results, however, it does appear that internal control had to be relied on more heavily in the Time task, probably because the passage of time did not itself provide a salient enough event to cue remembering. Thus, its influence could be seen more clearly in the time task data. Whatever the exact nature of these two classes of mechanisms, it is likely that they interact. For example, it might be that internal control helps regulate one's sensitivity to external cues—an increased level of self-initiated control might mean one is more susceptible to external reminders.

With regard to the finding that being physically in transition from one location to another is associated with increased thoughts about intentions, there are at least two possible explanations. First, it may be that transitional places are associated with lower cognitive load then places where people settle: people tend to be in places of transition when tasks have been completed or when they are taking a break from them. When in places such as a conference room or in an office, one is much more likely to be engaged in a task, and thus be mentally otherwise occupied. Second, location itself may also exert some influence over the kinds of thoughts one is likely to have. In other words, an office may have specific task associations (whatever task is currently "on the desk", for example), whereas being in places with no specific task associations may "free up" one's mind to think of other things, or to be more susceptible to the influence of external triggers. In either case, this finding points to an interesting line of research.

CONCLUSION

This experiment has begun to sketch out a general conceptual framework within which to understand the mechanisms responsible for thinking about and remembering intentions. What makes it distinct is that, in addition to being close analogues of tasks people generally cite as examples of prospective memory, these tasks were performed in a normal work situation, rather than in a laboratory. In this sense they were realistic in that they had to be performed despite a myriad of competing activities. In addition, the length of time over which they took place very likely means that the underlying mechanisms are different from those that operate in the short-term laboratory tasks.

It is difficult to imagine how some of the variables that might be of relevance to prospective remembering could be simulated in the laboratory. The discovery of the influence of factors such as physical location on thoughts about intentions provides an example. This is not to say that laboratory studies will not have a crucial role to play in contributing to our understanding of prospective memory, only that such an approach must necessarily constrain the variables under consideration. Given that this area of research is in its infancy, it is important that we do not stress experimental control to the detriment of exploration of those variables. We view this experiment as illustrative of how such exploration might be accomplished with the help of new technology.

Manuscript received 18 January 1995
Manuscript accepted 8 August 1996

REFERENCES

Craik, F.I.M. (1986). A functional account of age differences in memory. In F. Klix & H. Hagendorf (Eds.), *Human memory and cognitive capabilities: Mechanisms and performances.* (pp. 409–422). Amsterdam: Elsevier Science Publishers.

Edgington. E.S. (1972). An additive method for combining probability values from independent experiments. *Journal of Psychology*, *80*, 351–363.

Einstein, G.O., & McDaniel, M.A. (1990). Normal aging and prospective memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *16*(4), 717–726.

Eldridge, M., Sellen, A.J., & Bekerian, D. (1992). *The range, frequency, and severity of memory problems at work.* EuroPARC Technical Report No. EPC-92-129, Cambridge, UK.

Ellis, J.A., & Nimmo-Smith, I. (1993). Recollecting naturally occurring intentions: A study of cognitive and affective factors. *Memory*, *1*, 107–126.

Harris, H.E. (1984). Remembering to do things: A forgotten topic. In J.E. Harris & P.E. Morris (Eds.), *Everyday memory, actions, and absent-mindedness.* New York: Academic Press.

Harris, J.E., & Wilkins, A.J. (1982). Remembering to do things: A theoretical framework and an illustrative experiment. *Human Learning, 1*, 123–126.

Kidd, A. (1994). The marks are on the knowledge worker. *Proceedings of ACM Conference on Human Factors in Computing Systems (CHI '94).* 24–28 April, Boston (pp. 186–191). New York: ACM Press.

Kvavilashvili, L. (1987). Remembering intention as a distinct form of memory. *British Journal of Psychology, 78*, 507–518.

Kvavilashvili, L. (1992). Remembering intentions: A critical review of existing experimental paradigms. *Applied Cognitive Psychology, 6*, 507–524.

Lamming, M., & Newman, W. (1992). Activity-based information retrieval: Technology in support of personal memory. In F.H. Vogt (ed.), *Information Processing '92: Proceedings of the 12th World Computer Congress, Vol. III* (pp. 68–81). Amsterdam: Elsevier Science Publishers.

Lamming, M., Brown, P., Carter, K., Eldridge, M., Flynn, M., Louie, G., Robinson, P., & Sellen, A.J. (1994). The design of a human memory prosthesis. *Computer Journal, 37*(3), 153–163.

McDaniel, M.A. & Einstein, G.O. (1993). The importance of cue familiarity and cue distinctiveness in prospective memory. *Memory, 1*(1), 23–41.

Terry, W.S. (1988). Everyday forgetting: Data from a diary study. *Psychological Reports, 62*, 299–303.

Toothaker, L. (1991). *Multiple comparisons for researchers.* Newbury Park, CA: Sage Publications.

Want, R., Hopper, A., Falcao, V., & Gibbons, J.J. (1992). The Active Badge location system. *ACM Transactions on Information Systems, 10*(1), 91–102.

Wilkins, A.J., & Baddeley, A.D. (1978). Remembering to recall in everyday life: An approach to absent-mindedness. In M.M. Gruneberg, P.E. Morris, & R.N. Sykes (Eds.), *Practical aspects of memory.* London: Academic Press.