End-users searching the online catalogue

Kiestra, M.D.; Stokmans, M.J.W.; Kamphuis, M.

Published in:
The Electronic Library

Publication date:
1994

Article

End-users searching the online catalogue: the influence of domain and system knowledge on search patterns

M.D. Kiestra, M.J.W. Stokmans and J. Kamphuis

Tilburg University, Department of Language and Literature, PO Box 90153, 5000 LE Tilburg, The Netherlands.
E-mail: maaike.kiestra@pica.nl

Abstract: *In order to test the impact of system and domain knowledge on search behaviour in an online catalogue, an experiment was set up in a university library where students from three specialisation areas performed a number of search tasks in the online catalogue. The subjects differed in the amount of domain and system knowledge. In two sessions the subjects performed searches inside and outside their 'own' domain. During the first session all subjects had little system knowledge. After the first session, half of the group received instruction in catalogue use and the other half did not. To observe whether the induced differences in system knowledge had effects on the search performance, a second session was carried out. Subjects' search behaviour was videotaped and their comments recorded (they were encouraged to think aloud).*
Results show that the amount of system knowledge had a significant effect on search time as well as on the number of search patterns observed. Regarding domain knowledge, only one out of the six analyses concerning search time or the amount of patterns yielded a significant effect. A possible explanation for this result could be the questionable validity of the criteria used to distinguish between known and unknown domains. The difference in knowledge regarding familiar and unfamiliar domains is not as large as had been expected. The notion of end-users displaying habitual modes of behaviour is given considerable support by the data. This is reflected by the limited number of patterns observed.

1. Introduction

The development and use of electronic information retrieval systems in libraries has proliferated enormously in recent years. However, it is often noted that the design of these systems is not as user friendly as it could be (e.g. Baker & Lancaster 1991). This relies on the fact that it is often the designer in cooperation with library experts who designs the system. This leads to the situation in which the design of user-oriented interfaces is mainly based on implicit ideas of human information processing and on the experiences of the designers themselves, instead of taking the search behaviour of the end-users as a starting point. This method of designing interfaces possibly is caused by the difficulties encountered in circumscribing patterns in the search behaviour of end-users, which is often viewed as a continuing process of trial and error, as Bosman *et al.* (1992) already noted. The latter sounds like a contradiction to the often-heard notion of human beings depending heavily on habitual behaviour, saving time and effort by relying on routine and recurring modes of behaviour. These habits may express themselves in a limited number of search patterns.

The point made here is that user-friendly interfaces can be improved by facilitating the search routines provided by the system. Therefore it is necessary to uncover search patterns actually used by end-users and to identify possibly intervening variables.

From a consumer research point of view, searching a catalogue can be compared to the information acquisition phase within the consumer decision-making process (for an elaborate description of the consumer decision-making process, see for example Bettman (1979)). During the process of searching for information and deciding whether the information found is useful or not, different kinds of knowledge are used; among them are knowledge of facts (declarative knowledge) and knowledge of procedures (Anderson 1983). Declarative (what/who) knowledge may comprise names of authors worth searching for, and procedural (how) knowledge regards how to use these names in the catalogue to find titles. It is difficult to separate the effects of declarative and procedural knowledge; probably the two kinds of knowledge are interrelated.

Knowledge is viewed as a key variable in information searching processes (Allen 1991a). Variables often mentioned are knowledge of the retrieval system (system knowledge) and knowledge of the domain (domain knowledge). System knowledge changes, for example, by using different online catalogue systems or by following an instruction. Domain knowledge comprises the subject matter one wishes to acquire information about (Allen 1991a) and therefore varies with the subject for which information is searched. The assumption made here is that both domain and system knowledge are comprised of declarative and procedural knowledge. This could be an indication of a possible interrelationship be-

tween system and domain knowledge. Earlier research already pointed at the possibility of this interrelationship (Allen 1991b; Dimitroff 1992). One can reason that the system on which one is operating determines first of all which part of the system knowledge is activated but also which part of the domain knowledge will be useful to find a title. This evokes the research question of this paper, which can briefly be described as to what extent and in what direction search behaviour is determined by the amount of system and domain knowledge.

2. Method

This study focuses on the differences in the kind of search patterns used and the frequency in which these patterns are observed resulting from variations in domain and system knowledge. One of the methods used to answer research questions of this kind is the method of controlled task simulation as described in Lancaster (1993). Simulation studies are especially recommended because they enable researchers to combine the benefits of unobtrusive observation with experimental rigour.

2. 1. Design

In this study three variables are of interest: domain knowledge and system knowledge, both of which are independent variables influencing the third, dependent variable, which is the search behaviour displayed in catalogue use. These variables will be discussed successively.

Usually domain knowledge is defined as knowledge of the subject area (Allen 1991a). The amount of domain knowledge varies considerably between end-users (given a specific subject area). For any particular end-user the amount of knowledge differs per subject area. Given that individuals differ in knowledge regarding a specific domain, it is impossible to formulate in advance search tasks which are supposed to be on a known/unknown domain if the respondents participating in this study are selected on the basis of a random sample. Therefore it was decided for this study to approach groups of students specialising within the same study and having followed the same introductory courses. By choosing students who had followed the same introductory courses, the difference in amount of domain knowledge between individuals is probably not so large. It was assumed, however, that all students should be able to make at least some associations concerning all search tasks; it seems unrealistic to present a student in the humanities with tasks concerning chemistry.

The search tasks covered subjects in three specialisation areas, one of which was the domain in which the student was specialising. All subjects received the same search tasks (in and outside their domain). Consequently, a specific search task represents a familiar domain for one student and an unfamiliar domain for another student. Before the search session started, the subjects had to fill up a domain knowledge test in order to measure the amount of knowledge concerning the three domains (see Appendix). This test consisted of 10 true/false statements per domain. The subjects were asked to indicate the degree to which they agreed with the statement (7-point scale). In order to determine the amount of knowledge that a subject possessed regarding the domain, scores on these 10 questions were summed and divided by the total number of questions answered. So, a score of 7 is the maximum score one was able to attain on the domain knowledge test.

In this study system knowledge is defined as: 'knowledge of the ordering and categorisation principles of a catalogue, the references in a catalogue and the corresponding commands' (Kiestra & Stokmans 1994). In order to gain insight into the effect that the amount of system knowledge has on the kinds of search patterns used and the frequency with which these patterns are observed, the amount of system knowledge was manipulated such that one group of subjects possessed considerably more system knowledge than the other. In order to increase the amount of system knowledge some subjects received a thorough catalogue instruction.

Following each search session the subjects had to fill up a system knowledge test. This test was much like the domain knowledge test mentioned earlier. It consisted of 20 true/false statements on which the subjects could again indicate the degree to which they agreed with the statement.

At first sight, search behaviour can be studied by examining the separate actions performed by the end-user in a search session. An action can be defined as a (confirmed) 'input' to the 'request for information' by the system. If these separate actions are examined more carefully, one would rapidly conclude that one action in isolation has no meaning for the system or for the end-user. Consequently search patterns should be defined as a specific sequence of a number of *n* actions (*n* is a constant within an analysis but can vary between analyses), which is meaningful and manageable in the context of search behaviour in a catalogue. In our opinion, a meaningful unit should reflect an interaction between the system and the end-user, which can be interpreted uniquely. In many online catalogues a meaningful sequence of actions could consist of the sequence: 'select a key' (e.g. author); 'enter text' (name of the author); and the overt reaction to information provided by the system (for example, select one title out of a list of titles). On the basis of meaningfulness the minimum sequence of actions in this study is three (see also Figure 1). For this research this was also the maximum number which is manageable, because the number of different search patterns which theoretically

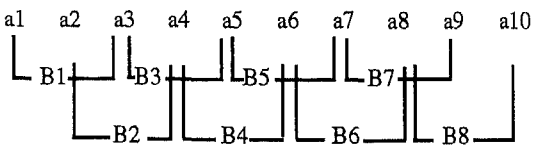

In which: ai (i = 1, 2, 3,...,10): the successive actions performed;
Bk (k = 1, 2, 3,...,8): the successive search patterns

Figure 1: *Defining search patterns*

can be defined in that case is extremely large. Consequently, the rather abstract term 'search behaviour' in this study is operationalised as a sequence of three separate actions (which in turn are defined as a confirmed 'input' to the 'request for information' by the system).

In the experiment each subject was scheduled for two search sessions, each consisting of four search tasks (two from a familiar domain and two from an unfamiliar domain). The order with which familiar and unfamiliar tasks were presented was randomised, as were the task contents.

Between the two sessions, half the subjects (randomly selected) received instruction (experiment groups) while the other subjects received no instruction (control groups). It was expected that domain knowledge will not increase between the two sessions, because no lectures were given in that period. System knowledge was expected to grow as a result of 'history', which means that subjects learn something about the system because of the search experiences they had during the first session, or other search sessions that they accomplished between the first and second sessions. This history effect will appear irrespective of intermediate instruction. The system knowledge will of course increase considerably for the experiment group after following instruction.

Consequently, eight groups of tasks are identified by varying systematically the system knowledge (instruction versus no instruction), the domain knowledge (familiar versus unfamiliar domain) and whether the search task regards the first or second search session.

The experiment can be classified as a pre-test-post-test control group design (Cook & Campbell 1979). This design has the advantage that differences found between the groups and between sessions can be attributed exclusively to differences in system and domain knowledge, because the subjects were randomly assigned to the control and experiment group and therefore all other variables that might be intervening (like age and level of education) are randomised.

2. 2. *Subjects*

Subjects in this study were 34 students of the Department of Language and Literature of Tilburg University. Of these subjects, five did not finish the entire experiment: they missed the post-test because of illness or lack of time. Therefore the results reported below are based on 29 subjects. In total they performed 232 tasks, of which 116 regarded an unfamiliar domain and 116 regarded a familiar domain.

After completing the experiment the subjects were given a small monetary reward.

2. 3. *Procedure*

The experiment took place in a secluded room at the library of Tilburg University, November/December 1993. The experiment was administered to each subject separately. The subjects were instructed to imagine a situation in which they had to prepare a paper on the given subject, for which at least three relevant titles had to be found. The titles found had to be written down. During the search session the subjects were allowed to use the written manual of the online catalogue and the UDC tables which were placed next to the computer terminal. The maximum amount of time allowed to accomplish one search task was 15 minutes. This was a pragmatic choice enabling a subject to perform four tasks within an hour. On the other hand, most search sessions generally last less than 15 minutes. This was checked by examining the connection times of Tilburg University Library OPAC over a one-year period. Consequently, search sessions lasted a maximum of one hour during which four search tasks had to be carried out.

Before starting with the first search task, the subject was asked to fill up the domain knowledge test mentioned earlier in order to measure the amount of knowledge concerning the three domains. Following the last (fourth) search task the subjects were invited to fill up the system knowledge test.

Two or three days after the first search session, half the subjects were given instruction regarding usage of the online catalogue. The other subjects did not receive instruction. After yet another two or three days the subjects participated in the second search session of the experiment. This second search session was designed in the same way as the first, except that this time other search tasks were used.

The search behaviour of both search sessions was videotaped and the comments of the subjects, who were asked to think aloud while performing the task, recorded. An observer who was monitoring the performance of the subjects on a additional screen was also present in the room.

3. Results

In order to check whether the manipulation of domain knowledge has succeeded, the average scores on the domain knowledge tests were compared between tasks concerning domains with which the subject was supposed to be familiar and those with which the subject should be unfamiliar. The average score on the task concerning domains with which the subjects were familiar was 4.07 (s.d. = 0.86). Average score on unfamiliar domain tasks was 2.02 (s.d. = 1.44). This difference is significant ($t = 13.17$, $df = 189.28$, $p < 0.05$). So, at first sight the manipulation of domain knowledge has succeeded; subjects had more knowledge regarding familiar domains than regarding unfamiliar ones. But the maximum possible score (7) is considerably higher than the average score regarding the familiar tasks. If the manipulation really would have succeeded, it means not only that subjects would know relatively more about a familiar domain but also that they would know much more in absolute terms. In order to test this premise, all tasks performed were divided into two groups — those which had domain knowledge scores below the median and those which had scores above it. This enabled the researchers to examine whether or not those tasks which scored below the median were indeed the tasks on which the subjects were supposed to have less domain knowledge (and the other way

Table 1: *Classification of the task on the basis of a priori familiarity and scores on the domain knowledge test.*

		domain *a priori*	
		familiar	*unfamiliar*
scores	high	91	25
	low	25	91

round). This cross-tabulation is presented in Table 1 (overleaf).

Table 1 shows that the *a priori* classification does not correspond with the classification on the basis of the domain knowledge test scores. Consequently the *a priori* classification, as well as the classification of the basis of the domain knowledge score, will be considered in all analyses when determining the effect of domain knowledge on the search behaviour of the respondent.

A: Domain knowledge operationalised as familiarity of the task

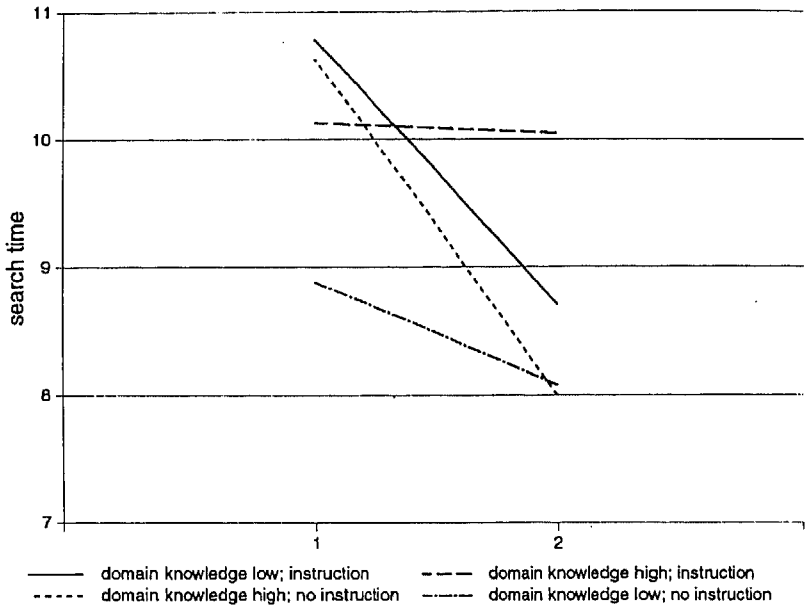

B: Domain knowledge operationalised as scores on the domain knowledge test

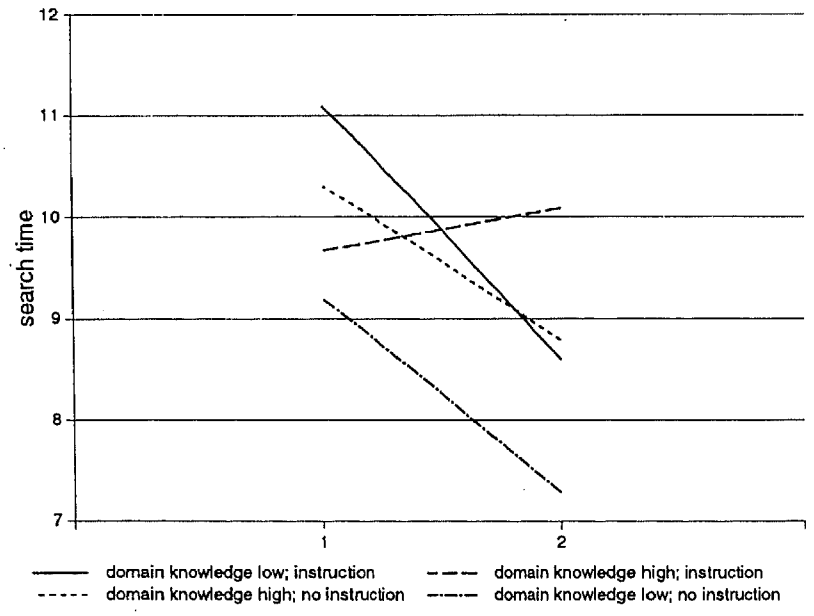

Figure 2: *Fluctuations in search time related to domain knowledge and system knowledge (increased by following instruction).*

3. 1. Time spent on the task

The maximum time one was allowed to spend on a task was 15 minutes. Seventy tasks (accounting for 30% of the total) were stopped because the time limit set had expired. The time spent on the task, once it was terminated by the respondent, was analysed regarding fluctuations in the time spent. This variable is normally distributed (Kolmogorov-Smirnov goodness of fit, $z = 1.16$, $p > 0.05$).

On average, subjects spent 9.35 minutes (s.d. = 3.09) on an search task. It was hypothesised that the time spent on a task depends on the familiarity of the domain, as well as the amount of system knowledge one possesses. Figures 2a and 2b show that the average time spent varies considerably between the eight groups of tasks identified in the design section.

Whether or not differences in the amount of domain and system knowledge (as well as whether it concerned a first or second session) had an effect on the time spent on a task was examined by means of three analyses of variance.* In the first analysis, domain knowledge was operationalised as the *a priori* familiarity of the tasks (based on the student's specialisation area). Here domain knowledge had no

* *Analysis of variance tests whether the groups distinguished (in our case the eight groups of tasks) differ on the dependent variable (in our case search time or amount of patterns). To test if the differences found in the dependent variable can be attributed to a particular variable or just random fluctuations, an F-distribution is used. This distribution depends on the degrees of freedom (v_1, v_2) in the testing. The degrees of freedom are indicated as subscripts of the F-value.*

significant effect on the time spent ($F_{1,157} = 1.530$, $p > 0.05$), while the difference in time searched between the first and second session, as well as the instruction variable, was significant ($F_{1,157} = 9.905$. $p < 0.05$ and $F_{1,157} = 5.445$, $p < 0.05$, respectively). In this analysis none of the interaction effects between domain knowledge, system knowledge and testing (first or second) were significant.

In the second analysis the operationalisation of domain knowledge regarded a categorisation of the tasks on the basis of the absolute score of the domain knowledge test (see also above). Now domain knowledge had a significant effect on the time spent ($F_{1,157} = 2.862$, $p = 0.09$). This effect is not very convincing, however. The difference in time searched between the first and second session, as well as whether or not one had received instruction, were significant ($F_{1,157} = 10.355$. $p < 0.05$ and $F_{1,157} = 5.766$, $p < 0.05$, respectively). In this analysis one of the interaction effects was significant, namely domain knowledge and session (first versus second) ($F_{1,157} = 3.022$, $p = 0.08$), but this effect is not very strong.

In the third analysis of variance, domain knowledge was treated as a covariate. Again similar results were found: the covariate domain knowledge did not yield significance ($F_{1,159} = 1.760$, $p > 0.05$), while session (first versus second) and the instruction variable had a significant effect on the amount of time spent searching ($F_{1,157} = 10.072$, $p < 0.05$ and $F_{1,157} = 5.125$, $p < 0.05$, respectively).

A: Domain knowledge operationalised as familiarity of the task

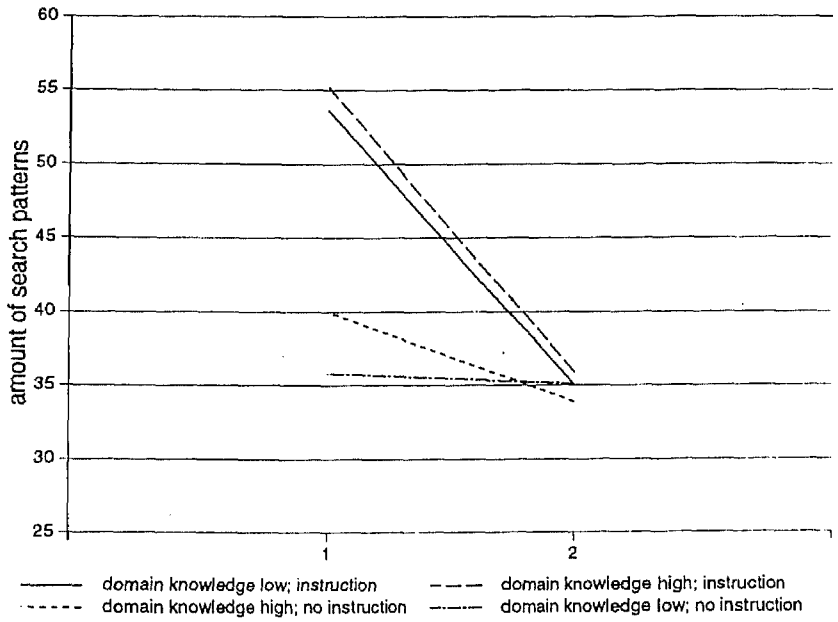

B: Domain knowledge operationalised as scores on the domain knowledge test

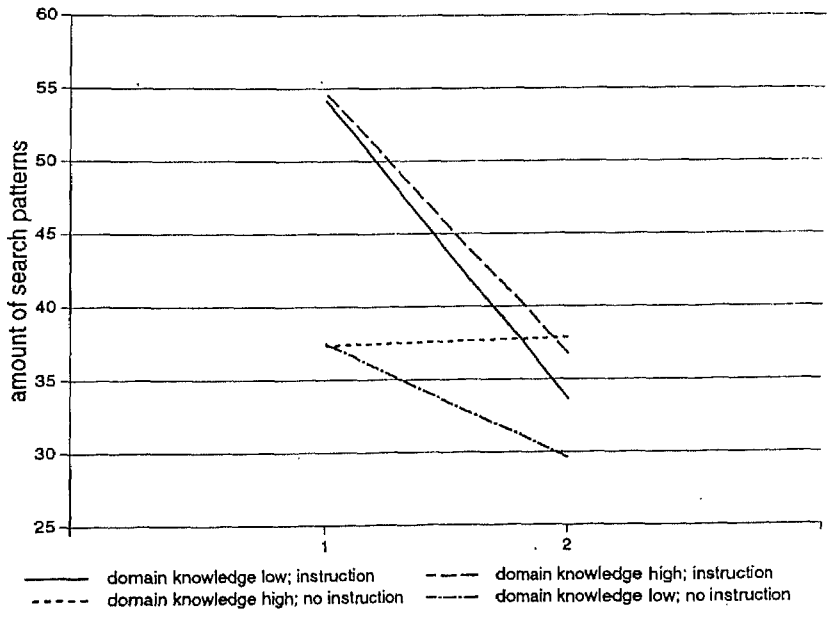

Figure 3: *Fluctuations in amount of patterns related to domain knowledge and system knowledge (increased by following instruction).*

3. 2. *Amount of search patterns*

Similar analyses were carried out regarding the amount of search patterns performed during a search task. Because of lack of time this analysis is based on 15 respondents only, comprising 120 search sessions. Figures 3a and 3b show that the average amount of search patterns performed varies considerably across the groups of search tasks distinguished in the design.

Again three analyses of variance were performed and the results were almost similar to those found regarding time spent. If domain knowledge is operationalised as the *a priori* familiarity, it has no significant effect on the amount of patterns exhibited ($F_{1,111} = 0.294$, $p > 0.05$), while session (first versus second) and whether or not instruction was received has a significant effect on the amount of patterns exhibited ($F_{1,111} = 10.090$, $p < 0.05$ and $F_{1,111} = 7.155$ $p < 0.05$, respectively). In this analysis only the interaction effect between session (first versus second) and whether or not one has received instruction was significant ($F_{1,111} = 5.206$, $p < 0.05$).

In the second analysis, the operationalisation of domain knowledge regarded a categorisation of the tasks on the basis of the absolute score of the domain knowledge test (see also above). This analysis generated results similar to the preceding

analysis on the amount of search patterns displayed. Again, domain knowledge had no significant effect on the amount of search patterns observed ($F_{1,111} = 0.733$, $p > 0.05$), while session (first versus second) and whether or not one had received instruction had a significant effect on the amount of patterns shown ($F_{1,111} = 10.082$, $p < 0.05$ and $F_{1,111} = 7.115$, $p < 0.05$, respectively). In this analysis again only the interaction effect between session (first versus second) and whether or not one had received instruction was significant ($F_{1,111} = 5.435$, $p < 0.05$).

In the third and final analysis of variance, domain knowledge was treated as a covariate. Similar results were found here as well. The covariate domain knowledge was not significant ($F_{1,113} = 1.222$, $p > 0.05$), while session (first versus second) and whether or not one had received instruction had a significant effect on the amount of patterns exhibited ($F_{1,113} = 10.235$, $p < 0.05$ and $F_{1,113} = 7.218$, $p < 0.05$, respectively). Again the interaction effect between session (first versus second) and whether one had received instruction was significant ($F_{1,113} = 5.383$, $p < 0.05$).

3. 3. Search patterns most frequently used

The total number of different possible actions to perform on the online catalogue system used is 42. This would indicate that the hypothetical total number of possible search patterns is 42*42*42= 74 008. The total number of patterns used in the 120 searches examined, however, was only 4725, divided over 553 different kinds of patterns. This result indicates the non-randomness of the way in which patterns are used. Analysis of the frequencies of these 553 patterns indicated that only 65 kinds of patterns were used on average by more than one subject.

The kinds of patterns most frequently used were determined by using three different criteria simultaneously. First of all, the pattern had to be observed more than 40 times, taking all search tasks together. Secondly, the difference in occurrence of the patterns between the first and second session had to be larger than 10. Finally, for each group (eight groups of tasks in total, as defined in the design section) the pattern should have been used at least once on average in the search tasks. These criteria resulted in 22 patterns. These patterns were also the most frequently used patterns presented at the Online Information Conference 1994 (Kiestra 1994), but these were based on the analysis of only 48 search tasks. This indicates that after analysis of all 232 search tasks, the patterns most frequently used will be these ones.

Hereafter, we will select three patterns to illustrate differences in search styles. The average frequency of occurrence of these patterns for each separate group of search task is shown in Table 2.

In the specific catalogue used in this study it is not possible to search on keywords. Possible search keys are: word(s) of the title, author/editor, corporation, words of congresses, person as subject, publisher, UDC, ISBN and ISSN.

Pattern 009 is a search pattern resulting in a list of titles, of which one is examined more closely. Regarding the subjects who had received instruction, this pattern is mainly used in the first session (before instruction) of the experiment and diminishes in the second session (after instruction). The non-instructed subjects displayed this pattern in the first as well as the second session regarding unfamiliar tasks. However, for familiar tasks this pattern vanished in the second session.

Pattern 015 is a search pattern indicating that a subject is scrolling in a list of short titles, or titles in full description. Regarding the instructed subjects, this pattern is mainly used in the first session (before instruction) of the experiment. Its frequency decreased in the second session (after instruction). For the non-instructed subjects the occurrence of this pattern also decreased, although not so dramatically.

Pattern 066 is a search pattern indicating that a subject is searching by means of UDC codes. The instructed subjects used this pattern only during the second session (after instruction) and not during the first session (before instruction). The subjects who did not receive instruction did not use this pattern at all.

Table 2: *Average frequencies of a search pattern per group of search tasks.*

Pattern (009)	*i/c*	*low sc. session 1*	*low sc. session 2*	*high sc. session 1*	*high sc. session 2*	*all 1*	*all 2*
1 select key: title word							
2 types a word	i	1.33	–	2.10	–	50	19
3 selects title full d.	c	1.70	1.23	1.20	–	42	32
Pattern (015)							
1 Enter: scrolling							
2 Enter: scrolling	i	19.25	2.25	14.13	2.80	457	72
3 Enter: scrolling	c	4.15	2.31	4.90	4.50	132	106
Pattern (066)							
1 select key: UDC							
2 types a code	i	–	1.50	–	1.00	6	34
3 chooses to limit	c	–	–	–	–	3	2

In which: i = instructed subjects; c = non-instructed subjects; low sc. = low score on domain knowledge test; high sc. = high score on domain knowledge test; - = mean frequency below 1.00 p/session; total frequencies per session in 'all' column.

4. Conclusion and discussion

The research question to be answered in this study regards the extent to which, and in what direction, search behaviour is determined by the amount of system and domain knowledge. In order to answer this question, end-users were observed while performing different search tasks. Time spent

searching for each task, and the amount and kind of search patterns used were examined.

The notion of end-users displaying habitual modes of behaviour is given considerable support by the data. This is reflected by the limited number of patterns observed (553), taking the possible number of patterns into account (74 008). But 553 different patterns used is still quite a lot. This large number of patterns displayed hints at the lack of restrictions imposed by the system, giving way to unbridled search behaviour. One can question the effectiveness of this kind of search behaviour. It is furthermore illustrated by the fact that of these 553 patterns, only 65 patterns are used by more than one subject, indicating that the search behaviour exhibited is very idiosyncratic. When analysing search patterns one should realise this and only analyse those patterns which are used by more than one subject.

The results furthermore showed that system knowledge had a significant effect on search time as well as on the amount of search patterns exhibited. Subjects who received instruction searched longer (especially in a familiar domain) but used less search patterns. This result may seem contradictory at first: one should realise, however, that subjects who received instruction used the UDC tables more frequently than subjects not having had instruction. When using these tables, one cannot search in the catalogue at the same time. So, the results obtained show an effect of system knowledge on search behaviour. A given catalogue, however, is designed for use by end-users with varying levels of domain and system knowledge. It is our impression, as well as Baker & Lancaster's (1991), that the system used would gain in flexibility and user-friendliness if the end-user was given the choice to do basic search (with only the elementary options: subject words, exact title words and author), or to perform elaborate search using more complicated options. This is suggested because we observed that the 'overload' of options that was available to some of the less experienced subjects in this experiment was confusing rather than facilitating their search, as can be concluded from the enormous amount of 'just try this one' patterns. Very few subjects know, for example, the meaning of 'Corporation', 'ISSN' or 'UDC'. The suggestion is to leave options like these out of the basis menu and class them as 'elaborate search commands'.

Regarding domain knowledge, the results indicated that this kind of knowledge had a significant effect in only one of the three analyses concerning search time. As regards the amount of patterns, no significant effect was found for domain knowledge. This does not necessarily mean that domain knowledge has no effect on search behaviour. The selection of tasks on which the subject should either possess much or little domain knowledge was not completely successful. The domain knowledge scores indicated that students were not as informed on a familiar domain as the researchers had hoped, and were more knowledgeable on domains with which they theoretically should be unfamiliar. The fact that most subjects were more knowledgeable on a supposed unfamiliar domain than the researchers had expected can be explained from the fact that all subjects had followed the same introductory courses, in which among other things the specialisation areas are reviewed in brief.

To summarise, the validity of the criteria used to distinguish between known and unknown domains can be questioned. As indicated above, the difference in knowledge regarding familiar and unfamiliar domains is not as large as the researchers had expected. Regarding the classification of subjects based on the domain knowledge test, one should realise that the domain knowledge test discriminates the subjects with above average knowledge from those below average knowledge. This does not mean, however, that the latter group did not know anything about the respective domains. Further research could elaborate on this issue by taking domains on more diverse areas of interest.

The results regarding search time are only based on those tasks which were accomplished within the time limit of fifteen minutes. Thirty per cent of the tasks were not accomplished within this time constraint. This was not expected by the researchers because close examination of connection times with the same catalogue indicated that most search tasks were finished within fifteen minutes. A possible explanation for this phenomenon could lie in the fact that the subjects were well (perhaps too well) aware of the experimental situation they were in and were very motivated to solve the search task. Some support for this suggestion can be found if the search times regarding the second search session are examined. These search times are considerably lower that those of the first search session (except for those of the group having received instruction and solving a search task on a familiar domain). If most subjects did not search the second session more effectively than they did the first session (which is very likely to be the case for the group that did not receive instruction), this suggests that subjects were more quickly satisfied with some of the titles that the system came up with in the second session than in the first session.

The question of whether the system operated on poses constraints upon the factual domain knowledge of the end-user could perhaps be answered more elaborately when all sessions are analysed. In addition a similar experiment would gain in depth if repeated using other catalogues with different designs.

References

ALLEN, B.L. (1991a) Cognitive research in information science: implications for design, *Annual Review of Information Science and Technology*, **26**, 3–37.

ALLEN, B.L. (1991b) Topic knowledge and online catalog search formulation, *Library Quarterly*, **61**(2), 188–213.

ANDERSON, J.R. (1983) *The Architecture of Cognition*, Harvard University Press, Cambridge, MA.

BAKER, S.L. and F.W. LANCASTER (1991) *The Measurement and Evaluation of Library Services*, Information Resources Press, Arlington.

BETTMAN, J.R. (1979) *An Information Processing Theory of Consumer Choice*, Addison-Wesley Publishing Company, London.

BOSMAN, F.E., M.A.H. ONGERING and H.P.A. SMIT (1992) Gebruikersvriendelijkheid van de pica online publiekscatalogus, Een evaluatieonderzoek met observaties van zoekgedrag, Pica internal paper, Pica, Leiden.

COOK, T.D. and D.T. CAMPBELL (1979) *Quasi Experimentation: Design and Analysis Issues for Field Settings*, Houghton Mifflin, Boston.

DIMITROFF, A. (1992) Mental models theory and search outcome in a bibliographic retrieval system, *Library and Information Science Research*, 14, 141–156.

KIESTRA, M.D. (1994). De invloed van domein- en systeemkennis op het zoekgedrag van gebruikers in de online catalogus van een wetenschappelijke bibliotheek, *Online Informatie Konferentie Nederland*, NBLC Uitgeverij, Den Haag, pp. 195–199.

KIESTRA, M.D. and M.J.W. STOKMANS (1994) What is system knowledge? A conceptualisation and measurement of 'system knowledge', Unpublished working paper.

LANCASTER, F.W. (1993) *If You Want To Evaluate Your Library . . .*, Library Association Publishing, London.

Appendix: Items used in the domain knowledge tests

First domain

(1) When children mature, watching television increasingly becomes a family activity while reading becomes an individual activity. Taken into account the importance of parents in literary development, their absence will decrease children's literary development.
(2) The amount of scientific research dedicated to children storytelling is negatively related to the degree to which children and parents engage in storytelling.
(3) Because literary experiences gained in a family affect the ways consecutive generations appreciate literature, one can observe family traditions when it comes to reading. External influences (school) only play a minor role.
(4) Empirical studies have revealed that watching television negatively affects reading behaviour.
(5) Children's knowledge concerning the function of written text precedes knowledge concerning the text format.
(6) When schools adopt views on 'emerging literacy', this hardly ever affects the ways in which reading proficiency is taught in these schools.
(7) A view on emerging literacy based on a selected subset of well known 'classic' books is not the right view on which to base research concerning emerging literacy.
(8) The definition of parent reading to children as a shared activity of two or more individuals with a marked beginning and ending, and involving an interaction centred round a text, is not sufficient.
(9) Parents having a functional view on reading are more inclined to leave parts of the reading education to schools than parents having a formal view on reading.
(10) Affective factors are more important than any other factors in the development of literacy in young children.

Second domain

(1) The current research on language attitudes does not attach much importance to the social-psychological approach, as was common during the 1960s.
(2) The attitude as intervening variable is much more important within the positivist/behaviouristic approach than within a mentalistic model.
(3) On both ends of a semantic differential, we always find two contrasting semantic adjectives.
(4) In the social comparison theories, convergence will only take place if the costs made by a person speaking counterbalance the rewards.
(5) Integrative motivation will yield a greater chance on success than an instrumental motivation in acquiring a language.
(6) Subtractive bilinguality defines situations in which a second language is learned at the expense of losing one's mother tongue.
(7) Social categorisation and in-group favouritism are not very important in Tajfel's social identity theory.
(8) A Dutchman calling other Dutchmen stingy can be regarded as an example of projective auto-stereotyping.
(9) The minimal group paradigm has clearly shown that the rewards of the in-group are minimised in order to maximise out-group differences.
(10) Foreigners, not having knowledge about the Dutch language, will associate texts spoken at low pitch as stemming from a less intelligent person than Dutch people will.

Third domain

(1) The activities of the [Dutch literary foundation] can be classified as a form of art policy in a broad sense.
(2) The activities of the [Dutch art organisation] can be classified as a form of art policy in a narrow sense.
(3) Two instruments used in the policy of [art foundation] are awarding prizes and assigning scholarships.
(4) According to [Dutch literary organisation] a translator of literary works should be given at least Dfl. 40 000 annually.
(5) The Royal Dutch Publishers Organisation employs activities on four different areas of literature policy.
(6) The [Dutch art organisation] is concerned with literary authors only.
(7) A contract between an author and the publisher also settles questions regarding film rights, translation rights and copyrights.
(8) Literary authors do not have the right to decide who publishes their books.
(9) The heirs of a deceased author have the right to claim the profits made by the sales of the books of this author.
(10) The government is inclined to introduce a system regulating the payments made by libraries to publishers (and, eventually, to authors).

The authors

Maaike Kiestra

Maaike Kiestra is Research Associate at the Faculty of Arts (Department of Language and Literature) at Tilburg University, and works at the Department of User Support at Pica Centre for Library Automation, Leiden. She is currently working on her PhD concerning the effects of domain and system knowledge on end-user search behaviour in online catalogues. Within the framework of her PhD she is also a visiting researcher at the Faculty of Arts at the University of Amsterdam.

Mia Stokmans

Mia Stokmans is Associate Professor at the Faculty of Arts (Department of Language and Literature) at Tilburg University. She graduated in economic psychology from Tilburg University, and finished her PhD concerning relative importance of product attributes in 1991 at Delft University of Technology. Her main areas of interest are consumer behaviour regarding leisure activities and research methodology.

Jan Kamphuis

Jan Kamphuis is Research Associate at a Dutch county library organisation. He graduated in economic psychology from Tilburg University. His main interests are consumer research and data analysis.