Browsing in Hypertext: A Cognitive Study

Erran Carmel, Stephen Crawford. and Hsinchun Chen

Abstract—With the growth of hypertext and multimedia applications that support and encourage browsing it is time to take a penetrating look at browsing behavior. Several dimensions of browsing are examined, to find out: first, what is browsing and what cognitive processes are associated with it; second, is there a browsing strategy, and if so, are there any differences between how subject-area experts and novices browse; and finally, how can this knowledge be applied to improve the design of hypertext systems. Two groups of students, subject-area experts and novices, were studied while browsing a Macintosh Hypercard application on the subject of The Vietnam War. A protocol analysis technique was used to gather and analyze data. Components of the GOMS model were used to describe the goals, operators, methods, and selection rules observed. Three browsing strategies were identified: 1) *search-oriented browse,* scanning and reviewing information relevant to a fixed task, 2) *review-browse,* scanning and reviewing interesting information in the presence of transient browse goals that represent changing tasks, and 3) *scan-browse,* scanning for interesting information (without review). Most subjects primarily used review-browse interspersed with search-oriented browse. Within this strategy, comparisons between subject-area experts and novices revealed differences in tactics: experts browsed in more depth, seldom used referential links, selected different kinds of topics, and viewed information differently than did novices. Based on these findings, suggestions are made to hypertext developers.

I. INTRODUCTION

BROWSING is a very common exploration strategy. While it is often discussed, compared and contrasted with other strategies, browsing is an activity that is not clearly defined. The purpose of this study is to better understand browsing, the basic cognitive processes associated with it, the impact of subject-area knowledge on browsing strategy, and how this knowledge might be applied to improve the design of hypertext systems. This experiment is a controlled study of browsers using a Macintosh Hypercard stack on the subject of The Vietnam War. Subject-area experts and novices are observed to study a wide variety of browsing behavior within a fixed setting and system.

An impetus for much of today's research in information retrieval (IR) is the advent of hypertext and, more recently, multimedia applications called hypermedia [35]. These systems allow nonlinear, direct access to vast amounts of information on a great variety of subjects [9]. As Nielsen [35] describes it, "Hypertext presents several different options to the readers,

and the *individual* reader determines which of them to follow *at the time* of reading the text."

Browsing is an activity that hypertext users commonly engage in. In fact, according to Nielsen [35], "the true purpose of hypertext is to provide an open, exploratory [browsing] information space to the user." Hypertext supports and encourages browsing by providing links between keywords and topics that can be explored at will as the focus of exploration changes; hypertext makes users feel that they are in control, free to explore according to their own needs [36]. This characteristic of hypertext clearly differentiates it from previous IR systems. Marchionini and Shneiderman [31] state that most other IR systems were designed for professional intermediaries, not to support end users in their task of discovery. They were designed to perform with the classical IR model that requires, as Bates *[2]* put it, "a single query presented by the user, matched to the database contents, yielding a single output set."

Several design problems arise from hypertext's support of browsing. The system can potentially confuse and disorient its user, the *embedded digression problem;* and it can cause the user to spend a great deal of time while learning nothing specific, the *art museum phenomenon* [18]. Hypertext systems provide bookmarks, context maps, backtracking or history lists, and guided tours to help the user overcome such problems [35]. To simplify navigation, hypertext systems employ a mouse, cursor, or touch-screen to select an object [40]. This object is usually a keyword that presents an imbedded menu of some sort, or a button that activates a graphic, video, etc. Yet, it is important to remember that there is currently "sparse evidence for . . . proven principles to guide design" [31]; it is unclear which combination of features is important to a given user of a given system performing a given task [35].

In spite of the new technology that supports it, there has been relatively little research on the basic human cognitive processes involved in browsing. By studying the operational mechanics of browsing, and by observing the impact of subject-area knowledge on browsing strategy, contributions to the design of hypertext systems may be made. **As** Nielsen [35] writes, "many of the most interesting insights leading to improved hypertext user interfaces have come from more qualitative observational studies."

II. LITERATURE REVIEW

A literature review was conducted to define browsing, understand the factors that effect browsing, understand browsing's functions, and examine previous experimental research in browsing.

A. Definition of Browse

The following discussion attests to some controversy about the definition of browsing. Using Webster's, it is impossible to operationally distinguish browsing from searching. Webster's defines browse as follows.

> To look through an aggregate of things casually especially in search of something of interest. To look over casually; to skim through a book reading at random passages that catch the eye.

Webster's defines search as follows.

> To go over or **look** through (a place, area, etc.) carefully in order to find something missing or lost. To explore or examine in order to discover. To look at, read, or examine a record.

By its first and most conservative definition, search is motivated by a very specific information need—one that persists until the goal is either found or abandoned for some reason. The second definition of search is somewhat more liberal, and the goal requirement is relaxed to the point where the search can be simply for something of interest— similar to the first definition of browse. In its most relaxed definition, no goal-orientation is implied at all, similar to the second browse definition. These definitions overlap so that browsing cannot be distinguished from searching.

Marchionini and Shneiderman [31] define browse as "an exploratory, information seeking strategy that depends upon serendipity . . . especially appropriate for ill-defined problems and for exploring new task domains." Liebscher and Marchionini [28] describe browse as an "alternative to the complex Boolean search strategy" for pursuing a "relatively simple and broad" query. And Marchionini [29] states that "Browsing often connotes an informal search process characterized by the absence of planning" using "techniques ranging from random and informal to systematic and formal." Cove and Walsh [13] use a three stage model to describe browsing:

1) Search browsing; directed search; where the goal is known
2) General purpose browsing; consulting sources that have a high likelihood of items of interest
3) Serendipity browsing; purely random

They describe browsing as a process of "picking out bits and pieces . . . selecting worthwhile and useful information" where the initial search criteria are only partly defined; to meet a "felt information need or interest"; "the art of not knowing what one wants until one finds it"; "essentially visual." The domain of browsing considered here corresponds to Marchionini's, only the domain is partitioned by the type of goal involved. In search browsing the goal is known—specific. In general purpose browsing the goal is interesting—general. In serendipity browsing there is no goal in the knowledge domain being browsed.

Salomon's [43] model is similar to that of Cove and Walsh:

1) Goal-directed search
2) Goal-directed browsing: users discover a goal in the course of browsing
3) Casual browsing: users casually look through numerous cards

Salomon's goal-directed browsing differs from Cove and Walsh's general purpose browsing in that the former suggests a component of goal reformation. As Bates [2] puts it, "These searchers are not just modifying the search terms used in order to get a better match for a single query . . . the query itself is continually shifting."

Bates [2] does not propose a taxonomy of techniques. Rather, she points out the need for using various techniques as required to meet the changing conditions a searcher may face. As new information is obtained in the process of search, the goal changes. This is called an evolving search. Strategies are selected in response to these new conditions to pick up new bits of information, a process called berrypicking. She sees browsing as a semi-structured, goal-oriented activity, distinct from Boolean search. She refers to area scanning as "the quintessential form of browsing in manual environments . . . flipping the pages and reading a passage here and there" getting the "feel or character of the author and his or her approach." This perspective of browsing corresponds most closely to the second of the three stages in the aforementioned models.

It is apparent that the browsing strategy employed depends on the type of task involved. Marchionini [30] classifies tasks as closed and open. Closed tasks have a specific objective, e.g., find the first year speed skating was introduced into the United States. They are goals for which a specific answer can be sought and found. They are often decomposed into subgoals. Control passes to subgoals as they are stacked and popped from the stack to move toward a goal state. This type of goal is required in order to apply formal Boolean search strategies. It is also required in order to use search browsing or goal-directed search, the first stage of the above browsing models. Open tasks have a general objective, e.g., find information about women who have travelled in space. This type of goal is required in order to use goal-directed and casual browsing, the second and third stages of the above browsing models.

Several operators and methods of browse/search have been identified. Liebscher and Marchionini [28] refer to scanning output for information of interest, using recognition tactics to select information for more careful examination. Marchionini refers to the use of a scan and select technique using general terms to locate information. He also measures lookup and examine moves of searchers. Bates [3] identifies four search activities: 1) Move: A basic identifiable thought or action; 2) Tactic: A handful of moves made to further a search (e.g., search formulation), 3) Stratagem: a larger, more complex set of thoughts and/or actions (e.g., citation search and footnote chase), and 4) Strategy: a plan for an entire information search that may contain moves, tactics, and/or stratagems.

B. Factors that *Effect* Browsing

Browsing strategy is individualistic and is subject to several variables. Marchionini and Shneiderman [31] developed a five point framework to identify determinants of success in seeking information. This framework also serves to identify the variables that effect browsing:

1) Setting—The physical environment, time allotted, access cost, plus the "functional constraints" of motivation and purpose—all factors that enable the search task.
2) Search system—The database structure and human-

computer interface that structures the knowledge and defines how it is accessed.

3) *Tusk domain*—*The* body of subject-area knowledge encoded by the system developer.

4) User—The user's mental model and experience with the task, system and search skills determine the user's orientation toward the search task; these are types of expertise. A user with a good mental model of the task domain will browse differently than one who is trying to comprehend the structure of the knowledge.

5) *Outcomes* are the product and process obtained from the search.

Mental models and *semantic networks* are concepts frequently used to describe cognitive processes as users interact with systems [5]. Marchionini and Shneiderman [31] describe a mental model as "a cognitive representation of a problem situation or system." It represents the content, structure, and relationship of information that helps the user to understand, draw inferences, and respond to conditions. It can be modified with new information and "run" internally as a knowledge base to generate hypotheses and make decisions. Norman [37] further divides this term into mental and conceptual models. A mental model is what the user has in mind and is not observable; while a conceptual model is given to the user by an outside source (e.g., a textbook, or hypercard system). Mental and conceptual models can be visualized as semantic networks, a concept originated by Quillian [42] as a mechanism for encoding and associating the meaning of words. These networks are made up of nodes and links. *Nodes* are units of information about something (e.g., a topic, idea, theme, point of view). *Links* are the relations between nodes through which the user navigates the network. Different types of links can be identified by link function. For example in hypertext, according to Conklin [12], *keyword links* allow navigation from index to information nodes; and *referential links* allow navigation between related information nodes. DeRose [15] describes a detailed taxonomy of link types.

Browsers are limited in hypertext to the relations encoded into nodes and links by the system developers. Problems may arise when the developer's conceptual model of a subject is .significantly different from the browser's mental model [8]. If the links do not identify meaningful relationships between nodes, they become superfluous to the browser. Lacking a good frame of reference, the browser can become "lost in hyperspace" [20]. Links are supposed to present relations that help the browser to integrate, synthesize, and associate nodes in the domain. For this reason, linking requires expert knowledge in addition to indexing knowledge. Indexing knowledge is required to provide points of entry into the nodes of data [13]. Expert knowledge is required to relate these nodes.

C. Uses of Browse

It seems clear that, from an IR perspective, browse is used primarily as an alternative to formal (e.g., Boolean) search strategy. According to Marchionini [29] and Bates [2] this occurs when:

1) the searcher cannot or has not defined a real search objective, a goal specific enough for a Boolean search. This is why professional intermediaries spend much of their time interviewing users to clarify their needs,

2) it is easier to browse than search. More formally, this is called cognitive laziness. This may reflect a poor understanding of the structure of information. Also, browsing is a system that is inherently active — keeping the fingers and eyes moving may be preferable to planning,

3) the system or library encourages browsing. This is true of a hypertext system. This must also include systems that simply do not support Boolean search, and

4) the user has a long-term commitment to an area of research and hopes to benefit from extraneous learning (incidental learning) in that area.

It seems apparent that people also browse to learn, where learning is the central focus rather than "incidental." For example, specialists browse to keep their expertise current. Also, people browse for entertainment, e.g., finding a new book to read, window-shopping, and museum-going.

D. Experimental Research on Browsing

Nielsen's extensive survey of studies [35] covers most of the important literature on hypertext, and related topics in IR. These studies primarily investigate hypertext design concepts and mechanics of implementation (e.g., the impact of interface software, hypertext versus scrolling text, hypertext versus traditional computing systems, hypertext versus paper, subjective judgements about hypertext, searching and IR, the uses of hypertext in education, individual differences, conceptual models, logging user interactions, observing users, and interface refinement). Similarly, The ACM Hypertext Conferences' have focused primarily on hypertext design concepts and mechanics of implementation. In this and other research, relatively few have studied the impact of variables on browsing or searching behavior (c.f., [13], [16], [28], [30], [36], [41], [44]). The focus of these studies is typically on the *performance* of certain design features of hypertext or on the *performance* of browsing strategies. None studied browsing specifically as the central focus.

Marchionini and Shneiderman [31] suggest there are historical and technological reasons that may explain why browsing has not been examined closely:

> Research related to on-line searching has focused on systems that aid professional intermediaries in finding a small number of "hits" in a large collection of records (library card catalogs) . . . They carefully plan in advance, consult thesauri, and combine terms in systematic and precise steps . . . Search intermediaries rarely browse informally, because focusing on the goal yields efficient and cost-effective performance.

Bates [2] echoes this view and adds that there is still a "lingering tendency in information science to see browsing *in contrast* to directed searching."

In summary, Marchionini and Shneiderman [31] state that research needs for full text systems are still very broad:

---

[1] cf. *Proc. ACM Hypertext '89 Conf.* and *Proc. ACM Hypertext '87 Conf.*
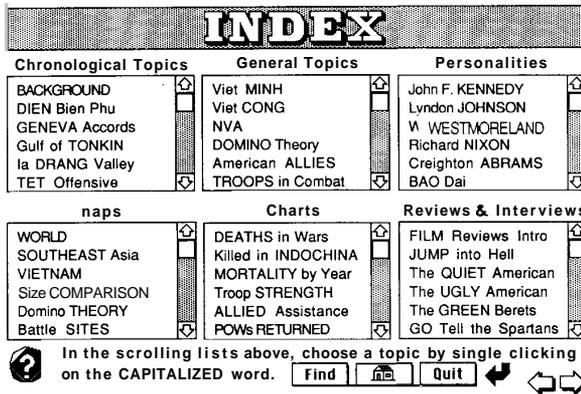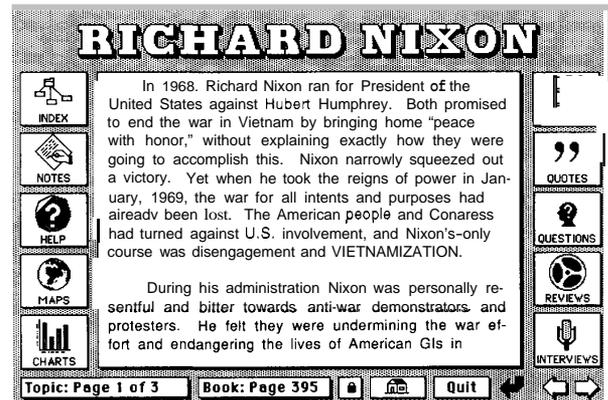
Fig. 1.   The Vietnam War Hypercard index.

Fig. 2.   The Vietnam War Hypercard topic.

We need to develop an understanding of the basic cognitive processes. We lack a clear understanding of the interactions of an information seeker's knowledge about the problem, past experience in searching for information and knowledge of possible sources of information.

Our research aims to address these issues by examining subjects' browsing strategies and associated cognitive processes in detail in a controlled, hypertext environment. We describe our research questions and experimental design in the following section.

III. RESEARCH QUESTION AND EXPERIMENTAL DESIGN

The browse-related research described in Section II suggested some general propositions for our research. They are as follows.

1) Numerous browsing strategies exist and can be operationally defined.
2) Subject-area experts browse differently than do novices.

The first proposition is suggested by the existence of the three stage browsing models proposed by Salomon [43] and by Cove and Walsh [13]. It should be possible to identify and classify basic operations or behaviors into procedures, and combinations of procedures into strategies. For example, Marchionini suggested some operations and phases such as scan and select, lookup and examine moves. Our experiment aimed to examine such browsing strategies in detail.

The second proposition is intended to study differences in browsing between subject-area experts and novices. Marchionini and Shneiderman [31] have identified this *task expertise* as a determinant of success in information seeking, as have Chen and Dhar [11]. Borgman [4] also cites research indicating that expertise in the task domain is a significant influence on search performance. It is not clear whether this influence extends to browsing, and exactly what this influence might be. Many of our dependent variables are aimed at measuring ways in which experts and novices could differ.

A. Experimental Design

The study employed a controlled experimental design. The setting was in a Macintosh computer laboratory for students at the University of Arizona. The search system used was

the Macintosh Hypercard system. Hypercard is capable of producing a large variety of menus and forms, is widely available, and is often used for developing and browsing hypertext [21]. The task domain was a Hypercard stack, a hypertext "computerbook" about The Vietnam War[2]. This stack was selected because it contained a combination of features likely to result in a variety of browsing behaviors and is a subject area of at least some interest to both expert and novice browsers. The stack presents an index (see Fig. 1), from which topics may be selected with an electronic mouse from one of several *categories — Chronological Topics, General Topics, Personalities, Maps, Charts, Reviews* and *Interviews*.

Selected topics are presented a page at a time (see Fig. 2). The user can page forward, backward, take an associative link (click on a capitalized word) to an imbedded menu in another topic, and return to previous topic or index. Most topics are between five and ten pages in length, There is a total of 455 cards. *Buttons* are also presented in several places to activate graphics such as maps and charts. *Special features* available from within each topic are presented in 1-in squares around the text from which the user may select 1) *Quotes,* which present quotes about the current topic, 2) *Questions,* which raise thought-provoking questions about the current topic, 3) *Sources,* a bibliography for the current topic, and 4) direct access to *Maps, Charts, Reviews,* and *Interviews.*

The subjects were scheduled for individual sessions in a computer lab. They were motivated to participate in the project by being offered extra credit in the course from which they were recruited. After brief introductions, the pretest was administered to classify subjects as expert or novice (described in detail in Section 111-B), the hypertext system was demonstrated to the subject, and the subject was introduced to the task. The task was for the subject to "look at a computer book about The Vietnam War" while thinking-aloud to "browse and learn whatever you like" for 30 minutes.

Data collection and analysis employed a verbal protocol analysis technique. This technique, explained in detail by Ericsson and Simon [17] involves recording a subject's actions,

[2]The Vietnam War, Version 2.0, Regeneration Software, 377 Collado Dr., Scotts Valley, CA 85066, Hypercard version 2.1 or newer

concurrent verbal report, and post-session interview, in order to understand a problem-solving process. A few minutes was spent with subjects to familiarize them with the concept of verbalizing their thoughts. Subjects were ensured that the experimenters had no personal stake in their observations about the stack, and that their honest feedback was most important. Most subjects needed at least a little prompting to start and restart the think-aloud process at points in the session. The following are representative of the type of prompting that was given:

Are you looking for something specific?
Why did you select that?
(Why) are you interested in X?
What was interesting about that?

It was occasionally necessary to give acknowledgment to the subjects to encourage their verbal feedback. These had to be free from directives and unbiased answers to opinions so as not to influence the subject's decision-making process.

At the end of the 30-min browsing session, a post-session interview was conducted to clarify points of strategy, special interests, etc. The following were standard questions asked of subjects during the interview.

Did you find this interesting?
Did you learn anything?
What was your strategy when you first sat down and confronted this?
Did you change your strategy over time? Why?

A videotape recorder was used to record the screens, keystrokes, the subject's concurrent verbal report and the post-session interview. The tapes were transcribed into problem-behavior graphs (PBG) [34], [45] to simplify analysis. PBG's describe problem-solving activities in a time sequence. Actions are recorded as *operator elements,* a finite set of actions that change the state of circumstances; knowledge is recorded as *knowledge elements* that specify what the subject knows about the task, the verbal protocol. Verbal protocols and PBG's are often used as a means to study problem solving processes [6], [11], [22].

Fig. **3,** at the end of Section IV-B, is an example of a PBG. The *Think-aloud & Commentary* section of the PBG is a transcript of the subjects think-aloud protocol during the session. Analysts' comments, bracketed, note nonverbal behaviors, e.g., a topic aborted or concluded, extensive scanning, and any assumptions made about behavior. The *Object* section presents the object of attention for the associated think-aloud. A simplified coding scheme developed by the authors was used to indicate these objects. For example, POWS, line 16, represents the topic *Prisoners of War*. The number that follows the dot represents a page number within the topic. The number 0 represents a keyword or referential link. For example, while browsing in JFK.l, line 21, the subject expresses an interest in *Special Forces* (SPEC.0), line 22, and selects it via the capitalized referential link. *An* additional alpha character may follow the number. For example, J means JUMP (a referential link taken), D means DOT (selected to activate a graphic display), and others indicate that special features were chosen. The *Operation* section codes the protocol, as described in detail in Section IV-A. The coded PBG was then further

analyzed to address the browsing strategies and expert-novice comparison propositions. We will examine this sample session more closely in Section IV.

B. Independent Variables

An adaptation of Marchionini and Shneiderman's[31] model of factors of success in information seeking was used as a framework to identify the independent variables that affect browsing:

External Variables

1) Setting (physical environment, time allotted, access cost).
2) Search System.
3) Task domain (conceptual model).

Internal Variables

4) Intelligence.
5) Motivation.
6) Purpose (task).
7) Task expertise (mental model & experience in task domain).
8) System expertise (mental model & experience with the system).
9) Research expertise (search skills).

The external variables were controlled as previously described in the experimental design. This was necessary in order to compare differences in browsing strategy between individuals rather than their setting or system.

Of the internal variables, task expertise was the only internal variable that was of interest to us in this research. Subjects were divided into two groups—experts and novices in the subject area of The Vietnam War. A pretest was developed for this purpose during the pilot study. Five subjects could not be classified due to limitations of the pretest. Of these, three that failed the test demonstrated expertise nonetheless from unusual military or other academic training that was not detected by the pretest. The remaining two passed the pretest due to good testing skills but demonstrated no expert knowledge of The Vietnam War while browsing. These five subjects were disqualified as experts or novices, and their results were excluded from statistical tests. Of the 31 subjects involved, ten were classified as experts. All ten of these were from The Vietnam War class. Sixteen of the subjects were classified as novices, 13 from the **MIS** class and **3** from the Vietnam class. The remaining five, previously described, could not be classified.

System expertise was discovered not to be a significant factor during our pilot test with eight subjects. Browsers adapted very quickly to the system. It was important that browsing behavior not be inhibited by the logistics of using the computer. Even first time computer users were able to navigate through the system easily within a few minutes and with minimal assistance.

Research expertise was a random variable. Students were selected as subjects from two comparable undergraduate classes—a political sciences course on The Vietnam War and

an introductory MIS class. None of them were, for example, research-oriented Ph.D. students.

Intelligence and motivation were also random variables. Motivation is "the goal or object of one's actions ... that which moves a person" (Webster's). Motivation is an internal response to the external environment. This study did not measure intelligence or motivation.

The purpose (task) given to the subjects, previously described, was an open one. The subjects were not given specific goals. They were free to browse, learn, or search anything and in any way they could. This study did not measure performance in completing any specific task.

C. Dependent Variables

Dependent variables were identified in this task environment in order to examine the browsing process and to identify differences in browsing behaviors between subject-area experts and novices. Classified by general area of interest, these twelve variables are defined below:

1) Browse Path: These variables help to understand the characteristics of paths taken while browsing.

    a) *Number of topics selected:* This is a count of the number of topics selected by the browser. Each topic is only counted once.

    b) *Referential Links:* Browsers may differ in their use of referential links. Referential links is a count of the number of topics selected from within a topic by clicking on a capitalized link word within the text. Browsers may tend toward a depth-first path with little use of these links, or toward a breadth-first path with intensive use of these links.

    c) *Topics aborted:* Browsers may differ in whether or not they complete their chosen paths. Topics aborted counts the number of topics not browsed to completion, through the last page of the topic. This includes topics that were never completed after referential links were taken, and topics that were abandoned part way through. This is indicative of a topic's attractiveness to the browser. Most topics are between five and ten pages in length. In Table I, topics aborted is measured proportionally to the total number of topics selected to normalize for differences in speed of browsing between experts and novices.

    d) *Special features:* Browsers may differ in their use of special features, the "gimmicks" of a well-designed hypertext. Special features counts the use of the links that surround the text area of topic screens: *sources, quotes, questions, maps, charts, reviews,* and *interviews.* These are, in effect, another way out of the current topic. Browsers may be attracted to them to escape the topic, for distraction, for their novelty, or in hopes of discovering additional information.

2) Search Behavior: These variables help to understand the subjects' tendencies to search. Browsers may vary on their tendency to search or to go without searching. This was suggested in the pilot study in which some experts immediately began to formulate and pursue specific goals (some of the pilot study experts were Ph.D. students with an unusual amount of research expertise). Search behaviors are characterized by the subject's spontaneous formulation of what we term a search goal. A search goal exists to accomplish a specific or closed [30] task; conversely, a browse goal has a general objective to accomplish a general, open [30] task.

    a) *Searches implemented:* This counts the number of times a subject attempted to find an answer to a search goal.

    b) *Searches not implemented:* This counts the number of times a subject articulated a search goal but did not attempt to find an answer.

3) Types of Topics Selected: These help to understand the tendency for browsers to select types of topics.

    a) *Unknown topics:* This counts the number of topics selected that the browser has never heard of before—selections based on negative recognition. Liebscher and Marchionini [28] indicated that recognition factors were important in topic selection.

    b) *Special interests:* Browsers may select topics that are related to their own areas of special interest. Examples are interests in U.S. Presidents, medicine, politics, and weapons.

    c) *Special interest topics:* This counts the number of topics selected that are related to the browser's special interests. as determined by the previous variable.

    d) *Expert topics:* This is a count of the number of topics selected that are based on expert knowledge of The Vietnam War, either because of existing expert knowledge or a desire to obtain expert knowledge. Browsers may select a topic that is very domain-specific (expert knowledge) or a topic that is very common sense. Common sense knowledge is "common knowledge about the world ... and the methods for making obvious inferences from this knowledge" [14]. The bulk of common sense knowledge is "superficial," "sensory," popular "consensus reality" [26]. In contrast, expert knowledge is acquired through specialized study, practice, and mastery over time [1].

The evaluation of an expert topic is based on the analysts' subjective assessment of the subject's selected topics. Examples of expert topics in the domain of The Vietnam War are Diem, Dien Ben Phu, Viet Minh, and the Tet Offensive. Examples of common sense topics are Nixon, drugs, and movies. Each of these had a role in The Vietnam War but are more typically related to common sense knowledge and special interests in the minds of novices. In Table I, expert topics was measured proportionally to the total number of topics selected to normalize for differences in speed of browsing between experts and novices.

4) Mental Model: These variables help to understand how information is perceived and used in the browser's mental model.

    a) *Integration type:* This is an evaluation of the primary use of information made by the subject. Different people use the same information in different ways and there may be patterns to this use. From the semantic network perspective of the mental model, one can either add nodes and links, change existing nodes and links, reinforce (confirm or acknowledge) existing nodes and links, or activate

(strengthen) existing nodes and links. Integration type was evaluated based on the browser's verbal protocol. The existence of an integration type of any kind indicates that the browser is using information from the conceptual model to influence his mental model in some way.

b) *Relating:* Relating counts the number of times that browsers compare their mental models to the hypertext author's conceptual model. This indicates that the browser perceives the conceptual model as something more than just a collection of information. For example, the statement "anything written from their side [liberal] is going to reflect . . . " indicates that the browser has positioned and labeled the conceptual model as a biased entity.

IV. PROCESS OF BROWSING BASED ON THE GOMS MODEL

This section reports results addressing proposition #1, which aims to define and differentiate browsing strategies. Browsing is clearly a cognitive information processing activity. Activities of this sort can be described using a GOMS model, introduced by Card, Moran, and Newell [7]. GOMS analysis describes a user's cognitive activity consisting of four components: 1) a set of *Goals,* 2) a set of *Operators,* **3)** a set of *Methods* for achieving the goals, and 4) a set of *Selection rules* for choosing among competing methods. The four components of GOMS correspond closely to the four components of the problem space *model — knowledge states* (initial, intermediate, and goal states), *operators* for changing states, *constraints* for applying operators, and *control knowledge* for deciding which operator to apply next.

GOMS models have been used to represent various perceptual, cognitive, and motor activities. They are commonly used to perform detailed keystroke, error, and performance time analyses [27], [38]. They can also provide approximation for quantitative and qualitative aspects of laboratory-based tasks [25]. For example, Card, Moran, and Newell [7] used GOMS to study text editing. They identified goals, e.g., to move to the next paragraph; operators to accomplish goals, e.g., push a certain key; alternative methods to accomplish goals, e.g., a series of operators to correct misspelling; and selection rules for different methods, e.g., use one method when a misspelling is on the same line as the cursor. In more recent research, Gong and Elkerton [19] used GOMS to study how novices learn new procedures in order to improve documentation. Their documentation was modified based on GOMS analysis of procedures. GOMS was used to compare keystroke and time differences between subjects. They found that the use of GOMS resulted in a significant improvement in user documentation as measured by the decrease in time required for novice users to learn new procedures. John [24] studied telephone toll and assistance operators (TAO) responding to customer requests. She used GOMS to measure and predict .performance times for operators completing various tasks in order to help specify design requirements for new workstations. After a decade of active research, GOMS models have become an effective research tool for both cognitive psychologists and systems designers (see Olson and Olson [39] for an excellent review).

Browsing seems to be amenable to description as an information-processing task using GOMS modeling. Specifically, we use the GOMS model to describe the qualitative and quantitative aspects of the browsing process. The following three subsections discuss the process of browsing in terms of the components of a GOMS model. Section V discusses browsing strategies as patterns of behavior in terms of this model.

A. *Goals*

Two types of search goals and two types of browse goals were distilled from the experiment. The search goal types were: 1) persistent search goals and 2) transient search goals. The browse goal types were: 3) transient browse goals, and 4) persistent browse goals. These are discussed in detail in the following sections.

1) Search Goals: Search goals have a specific objective to accomplish a closed task, mostly for fact-finding [30]. Examples of search goals observed:

- How many men were over there?
- Wasn't he the youngest president?
- Where were the troops located?

a) *Persistent search goals:* These are search goals that are the objects of a search — they are pursued (operated upon) until they are either satisfied or intentionally abandoned. Thirty-five different persistent search goals were identified by the number of *searches implemented* (dependent variable #5).

b) *Transient search goals:* These are search goals that are formulated but do not actually become the objects of a search. For example, they may have been evaluated as not worth finding, or just a wish for information. Twenty transient search goals were recorded from the protocol by counting the number of *searches not implemented* (dependent variable #6).

2) Browse Goals: Browse goals have a general objective to accomplish an open task [30]. With browse goals, one is not looking for a specific answer, but to gather information, change, reinforce or activate existing knowledge, depending upon motivation. The browse goal is a general goal that supplies some context-sensitive reason for the browser to gravitate toward the topic of interest. Sometimes the browse goal was simply to review something interesting, e.g., "because it looks interesting." Sometimes the browse goal was better articulated, e.g., "I want to see what they had to do with the war."

a) *Persistent browse goals:* These are browse goals that exist in the browsers' long-term memory. They are periodically activated in the browsing process as the user matches motivation, goals and information. The existence of persistent browse goals is inferred from the existence of long-term special interests demonstrated by browsers — patterns of topic selection and verbal feedback. For example, one novice subject demonstrated a special interest in people. Two-thirds of the topics she selected and much of her verbal feedback revolved around prisoners of war, drug addiction, race relations, war veterans, the war memorial, and a veteran's personal

account of the war. She completely avoided topics about the battles, politics, and weapons. In her words, "I wanted to find out more about people." The tendency to select topics of this type persisted over the course of her browsing, hence the inference of a persistent browse goal. In essence, motivation is operationalized in the browsing process as persistent browse goals.

b) *Transient browse goals:* These are browse goals that exist only temporarily, in short-term memory. They are normally replaced before being satisfied or intentionally abandoned. Examples of transient browse goals observed (italicized):

i) *I want to see SAIGON.* It was a major place in the war.

ii) This is interesting . . . *I want to see how these charts work.*

iii) *I'm going to find out more about TROOPS.* I was in ROTC for a long time.

iv) *Let's see about the Domino Theory.*

Approximately 350 transient browse goals were identified in the protocol. The transient nature of these browse goals is perhaps the most interesting and significant operational difference between it and the search goal. Queries "evolve" *[2]* around the transient browse goal; while novices discover goals in the course of browsing [43]. The goal-replacement process is as central to browsing as the goal-decomposition process is to searching. Novice #16 is a good example of this. Armed with only common sense knowledge about Vietnam, demonstrating no tendency toward any special interest, she selected predominantly expert topics (14 expert topics out of 19 number of topics selected). She was unfamiliar with many of the terms in these topics and used 13 referential links to explore them. *She never returned to complete previously-selected topics, and so aborted 15 of them.* As she read one topic, she became interested in another, and so on through the entire session. This tendency is typical for those that use referential links— there is a high correlation between the number of referential links taken and the number of topics aborted for browsers that took more than two links. Transient browse goals are activated and replaced in the browsing process as the user evaluates the match between motivation, goals and information. A detailed discussion of this evaluation process is presented in the operator *evaluate,* as follows.

B. Ouerators

Operators, in the GOMS model, are [7]:

elementary perceptual, motor, or cognitive acts, whose execution is necessary to change any aspect of the user's mental state or to affect the task environment . . . The behavior of the user is ultimately recordable as a sequence of these operations.

Six distinct types of operators were observed in the protocol. These are 1) select, 2) read, 3) integrate, 4) evaluate, *5)* search goal operators (stack goal and pop goal), and 6) browse goal operators *(add goal,* remove goal, and replace goal). The first

two operators are motor operator and perceptual operators, respectively. The last four are cognitive operators.

1) *Select* is the motor operator for locating and accessing information. In this experiment, each page of a topic is a separate topic node. Each node is physically accessed by pointing and clicking the electronic mouse on a link of some kind, or by positioning the eyes to the expected location of information, e.g., at a new sentence or paragraph.

2) Read is considered to be an elementary perceptual operator used to provide input to the mental model for subsequent mental processing. The function is that of a natural language processor—to obtain the unambiguous "semantic-pragmatic" meaning of the words [10].

3) *Integrate* is the cognitive process of using available information to activate the mental model or to produce some effect upon its content or structure. The dependent variable integration type measures possible effects of browsing on the mental model: add to, change, reinforce, or activate nodes and links in the mental model. Learning can take place during integration— modifications to the mental model may register in long-term memory; or the integration may only exist in short-term memory. Most subjects appeared to favor one or two integration types. The most common types were *add* and *change. Reinforce* and *activate* were present, but not as common. Two subjects had none—they exhibited no evidence of integration at all. Integration type was inferred based on analysts' observations and protocol analysis. An example of comments typical of each type are:

a) *Add*
   - I'm trying to find out about . . .
   - I never knew that he served under General Patton.

b) *Change*
   - I thought we lost a lot more than that.

c) *Reinforce*
   - That's what I thought.
   - They should have something in here on

d) *Activate*
   - I remember reading about counterinsurgency.
   - I forgot the French were there first.

Comments that included any of these integration types were interpreted as evidence of integration— that the subject was using information from the conceptual model to influence the mental model in some way.

4) *Evaluate* is a cognitive process that matches information with goals. Information is evaluated to determine whether it is pertinent to current goals. Goals are generated and then evaluated to determine whether they are pertinent to current information, within the constraints of motivation. The process appears to be iterative. It iterates until the "best" match is found

between information constrained by availability and goals constrained by motivation.

Evaluation gravitates attention toward certain goals and information. This heuristic is characteristic of a problem solving strategy called hill-climbing. Hill-climbing is essentially a process of 1) generating or locating a set of possible solutions from some current state, 2) moving to the "best" alternative state, and *3)* repeating this process [23]. In IR, search is a hill-climbing process that attempts to find the "best" information, where "best" is measured in terms relevance. That is, relevance is the heuristic for determining what information is "best." For example, Chen and Dhar [11] describe a "retrieval by instantiation" search strategy (a type of relevance feedback) adopted by on-line catalog searchers. Maron *et al.* [32], use the common IR metrics of recall and precision to rank the relevance of data to searchers' information needs— recall is the ratio of the number of relevant items retrieved to the total number of relevant items; precision is the ratio of the number of relevant items retrieved to the total number of items retrieved. Browsers also hill-climb; however, the evaluation of "best" is measured in terms of "interestingness." Interestingness is the heuristic for determining what information is "best." The term interestingness is used to refer to how interesting something is to the browsers— how strongly it attracts their attention. Interestingness heuristics have long been recognized as an essential component of an advanced and sophisticated learning paradigm— "learning by exploration/discovery" [33]. Lenat's **AM** program [26] is one of the few **AI** systems that identifies important new concepts based on interestingness heuristics without human supervision.

In this experiment, the word interesting was recorded 80 times; other times, the expression of interest was couched in terms of an explanation of why something was interesting. Here, again, are three of the transient browse goal examples:

a) *I want to see SAZGON.* It was a major place in the war.
b) *This is interesting* . . . I want to see how these charts work.
c) I'm going to find out more about TROOPS. *I was in ROTC for a long time.*

The italicized portions indicate that evaluation is taking place, and that interestingness is the heuristic in use. In the first example, the subject is expressipg interest due to recognizing a "major place in the war." In the second example, the word interesting is used explicitly. In the third, the subject is expressing interest in a subject due to past involvement. Browsers normally use interestingness as a criteria for evaluation. In contrast, searchers normally look for that which is most relevant to the goal.

Interestingqess seems to depend on a known or potential linkage of topics to goals. Most of these linkages required that the topics involved be recognized. Most browsers selected topics they recognized to some extent (e.g., "I've heard of this before"), with positive recognition. Occasionally, they selected topics with negative recognition (e.g., "I've never heard of this before"). Only 24 of the 285 topics selected were based on negative recognition. In these exceptional cases, the interesting link is with a goal to investigate "topics I've never heard of before" or to "learn about something new"—goals directed specifically toward unknown topics.

Sometimes, in browsing, the rate of hill-climbing is great—discernible progress is made in relation to a particular summit or browse goal. The ability to discern special interests in browsing behavior is a good example of how this climbing appears; the browser's transient browse goals and selected information revolve around certain persistent browse goals. In terms of the three stage browsing models of Cove and Walsh and of Salomon, this probably accounts for the perception of goal-directed or general purpose browsing (stage two). Sometimes, the climbing appears flat—no discernible progress toward any particular goal is made. In terms of the three stage browsing models of the literature review, this probably accounts for the perception of serendipity and casual browsing (stage three). While it is tempting to call such behavior random, it is somewhat misleading to do so. The behavior appears to be random in that browsing occurs around random presentations of information and that progress toward particular goals may not be discernible; however, the decisions themselves are not randomly made. They are guided by the heuristic use of interestingness. It is useful to view browsing as a hill-climbing process where progress toward a particular summit is a function of the interestingness of the goal at the summit relative to that of new peaks that are constantly appearing and disappearing from view. From this perspective, stages two and three of the models in literature review represent a hill-climbing continuum based on the relative influence of persistent browse goals on the browsing task— from no influence to total influence.

It is interesting to note the similarity between the evaluation function and Bates' description of the classical model of information retrieval. In this model [2], the searcher seeks a match between the document (information) constrained by its representatioq (availability) and a query (goals) constrained by an information need (persistent browse goals/motivation). In the search model, the information need is fixed and the heuristic for evaluation is relevance; while in the browsing process, the information need is often transient and the heuristic for evaluation is interestingness.

5) *Search Goal Operators:* Any decision to modify goals in some way requires operators for doing so. The operators required to modify search goals are *stack goal* and ***pop goal.*** When a goal is decomposed into subgoals, these subgoals are said to be "stacked" on top of the parent goal. Goal stacks can only be accessed from the top

of the stack; control passes to subgoals, one at a time, until they are either solved or abandoned. Subgoals are then—popped—from the stack to return control to the parent goal. Goal reduction is often considered to be a general, rational human problem solving method [46]. This does not mean to imply that the cognitive search process is simply a first-in, last-out process. In reality, control passes between decomposition processes and other processes in complex and unpredictable ways.

6) *Browse Goal Operators:* The operators required to reset browse goals are *add goal, remove goal,* and *replace goal.* Browse goals are independent of each other, not stacked and decomposed like search goals for planned operation. They are independent in that they may be added and pursued at will and in any order. *Add goal* occurs when a goal is added or activated as a result of the evaluation process. *Remove goal* occurs when a browse goal is forgotten (deactivated), satisfied, or consciously abandoned. *Add goal* and *remove goal* usually operate simultaneously—as a new goal is added (or activated) the old one is removed (or deactivated). The replace goal operator can be used to refer to such an event.

Now that all the operators have been identified. it is possible to describe the PBG protocol in detail. Please refer to Fig. 3, a composite example representative of the 31 sessions. The initial motivation (to "browse and learn whatever you like") was represented by the browse goal in line 1. The subject selected the index node, read a portion of the node, and evaluated it for interestingness. This cycle was repeated with insignificant modification of the goal until the subject evaluated a keyword as potentially interesting. This resulted in line 2, a replacement of the browse goal with another, "to see how the war started." A keyword link was selected to move to the first page of the topic BACKGROUND. **A** portion of this node was read and integrated, as evidenced by line 3, "I forgot the French were there first." This appears to be of integration type *activate.* The subject has used the information to remember a fact. The subject has also evaluated this fact as "interesting" (line 3). **As** a result of this evaluation, the subject continues in the current topic with insignificant modification of goals, selecting and reading page *2.* Additional integration took place in line 4, "The U.S. intervened . . . " This appears to be of integration type *add.* The subject has activated a portion of his mental model, added a node for Russia (sic), and added new links between the U.S., Russia, and The Vietnam War. Of course, whether this has been learned is unknown— learning was not measured. The integration may only exist in short-term memory. The subject continued this process, line *5,* by selecting, reading, integrating and evaluating. **At** this point, the evaluation resulted in a conclusion that this node was no longer interesting, "I know most of this stuff." The browse goal was nonverbally replaced as the subject directed attention to the *special features* surrounding the textual node, line **6.** The subject selected, read, evaluated, and replaced the browse goal with the new one, line 7, to browse movie reviews. Before selecting the link to movie reviews. the subject formulated and stacked a search goal, line 8, to find the movie reviews in the index. He selected the index, read and evaluated until he

found the relevant portion of the node containing the keywords that satisfied the search goal, and popped the goal. The browse goal was replaced on line 9, where the subject decided to look for realistic movies. He selected, read and evaluated them as uninteresting on line 10, with "I've seen most of these." This process continued until, on line 11, a movie was evaluated as potentially interesting because it was an exception, "one I haven't seen." The goal was replaced with "let's see if it's different." On line 12, the subject selected, read, and integrated information, "this is based on a diary . . . " On line 13, he evaluated this as an interesting exception to movies based on "the American point of view." The remainder of the example is similarly coded with operators.

Section IV-C simplifies the above description of the PBG protocol by aggregating operators into groups called *methods.*

C. Methods and Selection Rules

Methods and Selection Rules, the third and fourth parts of the GOMS model, are [7]

> learned procedures that the user already has at performance time; they are not plans that are created during a task performance. They constitute one of the major ways in which familiarity (skill) expresses itself.

It becomes apparent after coding operators for a short while, as in Fig. 3, that certain sequences of operators repeat under similar conditions. These sequences are referred to as methods or procedures. The conditions under which these methods are applied are called selection rules. Goals are achieved by applying selection rules to use the appropriate methods.

The browsing process observed in the protocol can be described as a procedure to satisfy goals by selecting one of two methods of behavior. These are referred to as 1) scan and 2) review. *Scan* is used to find (lookup) pertinent information. *Review* is used to integrate (examine) information. These are discussed in detail below. Methods and selection rules are summarized in **,** a browsing process model (BPM). A GOMS description summary of our research is presented in Fig. 4.

Scan: Scan is comprised of three operators: 1**)** select, 2) read, and 3) evaluate. Scan is used by the browser to lookup or find links either to: find other keywords for continued scanning, or to find information for review. In terms of its operators, scan is used to *select, read,* and *evaluate* the match between information and goals. For example, the PBG in aggregates all sequences of the operators *select, read,* and *evaluate* in brackets as scan methods. In each case, the intent of the sequence was to lookup, locate or find information that is pertinent to goals. On line 1 of the PBG, scanning is first used to find something interesting, a browse goal. It is used again for the same purpose on line 7 after the subject tired of reading background. Scan was used on line 8 in pursuit of a search goal to find certain keywords. On line 9, it was used to locate an interesting movie. On line 15, it was used to find an interesting chart. Scanning was consistently used in this way across all 31 subjects. Scan is listed, with its operators, on lines 19–22 of the BPM in Fig. 4.

Boolean search, consisting of query formulation and execution, may be considered an alternative to the scan method to find relevant information when keywords are known and a

THINK-ALOUD & COMMENTARY	OBJECT	OPERATION
1 Let's see if there's anything interesting here.	INDEX	Replace Browse Goal Scnn: {Select, Read, Evaluate}
2 Basically to see how the war started	BACKGR.0	Replace Browse Goal **Review:** {Select,
3 Interesting. I forgot the French were there first.	BACKGR.1	Read, Integrate, Evaluate, Select,
4 The U.S. intervened & Russia came in after we did.	BACKGR.2	Read, Integrate, Select,
5 I know most of this stuff. I learned this in class.	BACKGR.3	Read, Integrate, Evaluate)
6 [Directs attention to special features]	BACKGR.3S	Replace Browse Goal
7 This is an introduction to movies?	MOVIES.0S	Scan: {Select, Read, Evaluate} Replace Browse Goal
8 [Returns to index & searches for MOVIE keywords] [Search Goal found]	INDEX	Stack Search Goal Scan: {Select, Read, Evaluate} Pop Search Goal
9 Movies are exciting & glamorous accounts. I don't think they explain it all. I'm trying to find out something about that.	INDEX	Replace Browse Goal Scan: {Select, Read,
10 I've seen most of these movies.		Evaluate}
11 This is one I haven't seen... let's see if it's different.	MOVIEX.0	Replace Browse Goal
12 This is based on a diary of an unknown Viet Cong.	MOVIEX.1	**Review:** {Select, Read, Integrate,
13 Most are from the American point of view. That's very interesting. [End of topic]	MOVIEX.2	Evaluate)
14 I want to find something else interesting	INDEX	Replace Browse Goal
15 [Returns to index. Sees CHART keyword]	INDEX	Scan: {Select, Read, Evaluate}
16 This looks interesting. It reminds me of Rambo where he brings back those POWs.	POWS.0	Replace Browse Goal **Review:** {Select, Read, Integrate,
17 Is there any way to find out how many POWs were over there? [Search Goal abandoned without search]	5.POWS	Evaluate) Stack Search Goal Pop Search Goal
18 I want to find something new	INDEX	Replace Browse Goal
19 [Looks for something interesting]	INDEX	Scan: {Select, Read, Evaluate}
20 I want to know how much Kennedy was involved.	JFK.0	Replace Browse Goal **Review:** {Select,
21 16,000 advisors? Gee. That's a lot of advisors!	JFK.1	Read, Integrate, Evaluate}
22 The green berets. I'm interested in that.	SPEC.0J	Replace Browse Goal
23 I didn't know they were so involved over there	SPEC.1	**Review:** {Select, Read, Integrate ...}

Fig. 3.   Composite problem behavior graph coded with operators.

search mechanism is provided. For example, on line 8 of the PBG in Fig. 3, the scan for movie keywords could have been replaced with a query on the keyword MOVIE. The process of a Boolean search is represented by lines 16–18 of the BPM in Fig. 4. The Boolean search process was placed within the model based on our understanding of the potential of Boolean search from secondary literature review. The system did not provide the opportunity for subjects to use Boolean search.

Review: Review is comprised of four operators: 1) select, 2) read, 3) integrate, and 4) evaluate. Review is used to *integrate* pertinent information into the browser's mental model. In terms of its operators, review is used to select, read, and integrate information from nodes that have been evaluated as being of potential pertinence to goals. For example, the PBG in Fig. 3 aggregates all sequences of the operators *select, read, integrate,* and *evaluate* in brackets as review methods. Review is used in the PBG on lines 3, 4, and 5 to integrate information from the topic BACKGROUND in response to a browse goal to "see how the war started." Review is used on lines 12 and 13 to integrate information from a movie review in response to a goal to "see if it's different." Review was consistently

used in this way across all protocols. Review is listed, with its operators, on lines 27–31 of the BPM in Fig. 4.

The distinction between the scan and review methods is of fundamental importance. The focus of attention during scan is on *evaluation;* the focus of attention during review is on *integration.*

Pursue–Goal: The decision to operate on goals is modeled in the *Pursue–Goal* procedure of the BPM in Fig. 4, lines 7–11. In order to be operated upon, goals must contain two parts: 1) an *intention* to either find or to integrate information, and 2) an *object* to either find or integrate, such as a specific fact or general information. When the intention is to find an object, either Boolean search or scan is selected. Boolean search is used when query formulation is desired and possible (line 9); otherwise scan is used (line 10). When the intention is to integrate information, review is selected (line 11).

Modify–Goal: The decision to modify goals is modeled in the *Modify–Goal* procedure of the BPM in Fig. 4, lines 12–14. This includes the *modify_search_goal* and *modify_browse_goal* procedures (lines 28–34). The decision to modify search goals differs from the decision to modify

```
 1  Browse (goal)
 2      Control strategy                          control strategy for EVALUATE operator
 3        Type: Hill-climbing
 4        Heuristic of pertinence:
 5          Relevance  ...                        if goal is search goal (specific, closed, or a subgoal)
 6          Interestingness  ...                  if goal is browse goal (general, open, not a subgoal)
 7      Pursue-Goal                               to Boolean search, scan, or review
 8        [select either ...                      to find information
 9          Boolean-Search (line 16) ...          if query desired and possible
10          Scan (line 19)]
11        Review (line 23) ...                    to integrate information]
12      Modify-Goal                               to modify goals
13        Modify-Search-Goal (line 28)           if search goal recourse is required
14        Modify_Browse_Goal (line 31)           if browse goal recourse is desired
15      Go to Browse (line 1)                     if goal exists

16 Boolean-Search
17    Formulate-Query
18    Execute-Query (procedure)

19 Scan
20    SELECT
21    READ
22    EVALUATE (pertinent match of information and goals)

23 Review
24    SELECT
25    READ
26    INTEGRATE (add to, change, reinforce, or activate mental model)
27    EVALUATE

28 Modify-Search-Goal
29    STACK GOAL ...                              if goal(s) or subgoals(s) required
30    POP GOAL...                                 if goal(s) satisfied or goal(s) abandoned

31 Modify-Browse-Goal
32    ADD GOAL ...                                if new goal(s) desired
33    REMOVE GOAL ...                             if goal(s) forgotten, satisfied, or abandoned
34    REPLACE GOAL ...                            if concurrent 'add goal' and 'remove goal'
```

Fig. **4.** Browsing process model.

browse goals. Search goals are only modified when there is no alternative or recourse. Recall that search goals are usually persistent —— perated upon until they are satisfied or consciously abandoned, and decomposed if necessary. When search goal modification is required, search goals are stacked and popped. The *modify_search_goal* procedure (line 13) enables the application of any combination of *stack goal* and *pop goal* operators to be applied (lines 29 and 30). The consideration of how to modify goals is a result of evaluation during scan or review. In contrast to search goals, browse goals are modified at will, as desired — recall that browse goals are transient, exist independently in short-term memory, and are normally replaced before being satisfied or intentionally abandoned. The *modify_browse_goal* procedure (line 14) enables the application of any combination of browse goal operators *add goal,* remove goal, or replace goal (lines 32–34). Again, the consideration of how to modify goals is a result of evaluation during scan or review.

Control Strategy: Lines 2–6 of the BPM in Fig. 4 identify the heuristic used for the control strategy. Recall that the *evaluate* operator, invoked during scan and review, iterates toward a match between information constrained by availability and goals constrained by motivation; and that this process of hill-climbing is based on a heuristic of pertinence to goals. For search goals, line 5, the measure of pertinence is *relevance.*

For browse goals, line 6, the measure is *interestingness.*

As previously indicated, the browsing process is iterative in nature. This is represented by *browse,* the controlling procedure of the BPM in Fig. 4, lines 1–15. This procedure repeats as goals are modified (*modify_goals*) and methods of behavior are selected to pursue them (*pursue_goals*). Line 1 indicates that browsing requires a goal. Line 15 indicates that browsing is iterative.

V. BROWSING STRATEGY

With the identification of goals (G), operators (O), methods (M) and selection rules (S) in the previous section, it is possible to address the first of our two research propositions: that numerous browsing strategies exist and can be operationally defined. Protocol analysis revealed three distinctive browsing strategies. We call them 1) scan-browse, 2) review-browse, and 3) search-oriented browse. In terms of the browsing process model (BPM) in Fig. 4, each strategy iterates through decisions that involve a different set of methods. Fig. **6** illustrates the methods involved in each of the three browsing strategies, which are described in detail below.

1) *Scan-Browse:* In scan-browse, the focus of attention is on scanning to evaluate information for interestingness. The browser cycles through decisions that involve the methods *modify_browse_goal* and *scan,* but do not involve review. The

Goals:
  Search *Goal:* specific goal, open task, or subgoal, fact-finding orientation
    Persistent: search operations performed on goal
    Transient: no operations performed on goal, e.g.: wish, justification
  Browse *Goal*: general, open, not a subgoal, independent, information-processing orientation
    Persistent: in long-term memory, effect of motivation
    Transient: in short-term memory, ever-changing, context-sensitive

Operators (to accomplish goals):
  SELECT: physical or mental activity of selecting a node, e.g. via link with a mouse or
    or via repositioning of eyes
  READ: natural language processing to obtain semantic-pragmatic meaning of words
  INTEGRATE: use of information in conceptual model to effect the content or structure of the
    mental model — add to, change, reinforce, activate nodes and links
  EVALUATE: iterative adjustment toward a match between goals constrained by motivation and
    information constrained by availability. Results in a decision to modify goals and
    to select nodes ⋅
  STACK *GOAL:* decompose goal and pass control to search subgoal(s)
  POP GOAL: delete search subgoal(s) solved or abandoned to return control
  ADD *GOAL:* generate or activate browse goal(s) independent of existing goal(s)
  REMOVE *GOAL:* forget, deactivate, or consciously abandon browse goal(s)
  REPLACE *GOAL:* Concurrent add goal and remove goal

Methods (procedures) and **Selection Rules** (conditions):
  Browse: the process of 1) applying *pursue_goal* to either lookup or integrate information, and
    2) applying *modify_goal*.
  Pursue_Goal: the process of choosing either scan or *Boolean_search* to find **(lookup)**
    information, or review to integrate information.
  Modify–Goal: the process of choosing to *modify_search_goal* and/or *modify_browse_goal*
    in response to evaluation during *scan* and review.
  Boolean–Search: the process of formulating and executing queries using keywords and operators
    such as AND, OR, and NOT to find information relevant to goals.
  Scan: the process of selecting, reading and evaluating information and goals with the intent of
    finding pertinent information and goals. Focus of attention is on evaluating.
  Review: the process of selecting, reading, integrating information expected *to* be pertinent *to*
    goals; evaluating information and goals for actual and potential pertinence. Focus of
    attention is on integrating.
  Modify–Search–Goal: the process of modifying search goals by applying any of the search goal
    operators stack goal and pop goal.
  Modify–Browse–Goal: the process of modifying browse goals by applying any of the browse goal
    operators add goal, *remove* goal or replace goal.
  Control Strategy: heuristic for evaluating pertinence of information and goals during
    evaluate: interestingness for browse goals, relevance for search goals.

  Note: See the Browsing *Process Model*, Figure **4**, for the complete specification of selection
    rules in the integrated GOMS model.

Fig. 5. GOMS description summary

browser scans for information of interest, but does not review to integrate the information found.

In terms of the BPM in Fig. 4, the browser begins with a browse goal, and therefore adopts the heuristic of interestingness for the control strategy (line 6). The browser enters *pursue_goal* (line 7) with the intention to find information (line 8), and therefore selects scan (line 10). Scanning (line 19), the browser selects, reads, and evaluates for a match between information and goals; for example, when scanning an index for something interesting. When no information worth integrating is found, the browser may enter *Modify_Goal* (line 12) to modify browse goals (line 14). Modifying browse goals (line 31), the browser adds, removes, and/or replaces goals; for example, to find something else interesting. This browsing cycle repeats as illustrated in Fig. 6.

Two subjects engaged in extensive scan-browsing421 (a novice) and #9 (not classified). Other subjects selected topics and then reviewed them; these two selected topics and then scanned them. Each aborted all topics with referential links to other topics, which were also scanned. For example, subject #21 started to "just check out" NIXON, quickly returning to the index. He was "kinda interested" in KENNEDY, quickly returning again. The subject was "a little interested in weapons" and paged through several screens "just scanning [his word]" to "see if anything sounded more interesting." Questions designed to elicit evidence of integration (cf., Section III-A) failed to find any support. This continued for 30 min.

2) Review-Browse: In review-browse, the focus of attention is usually on reviewing to integrate information into the browser's mental model, and browse goals are freely modified for this purpose. The browser cycles through decisions that involve the methods *modify_browse_goal*, *scan,* and *review*. As in scan-browse, the review-browser occasionally scans to find information. Twenty-nine of the 31 subjects review-browsed, based on evidence of information integration in the PBG's that indicated that review was taking place.

In terms of the BPM in Fig. 4, the browser begins with a browse goal, and therefore adopts the heuristic of interestingness for the control strategy (line 6). The browser enters
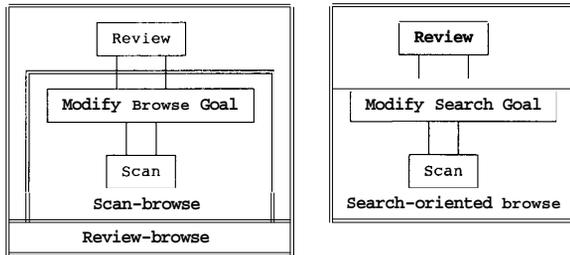
Fig. **6.** The three browsing strategies: scan-browse, review-browse, and search-oriented browse.

pursue_goal (line *7*) with the intention to find information (line 8), and therefore selects *scan* (line 10). Scanning (line 19), the browser selects, reads, and evaluates for a match between information and goals; for example, when scanning an index for something interesting. When the browser finds what appears to be interesting information, the browser enters *modify_goal* (line 12) to modify browse goals (line 14). Modifying browse goals (line 31), the browser adds, removes, and/or replaces goals in order to integrate information; for example, when deciding to explore a specific topic. The browser then reenters *pursue_goal* (line 7), this time with the intention to integrate information from the conceptual model with the browser's mental model, and therefore selects *review* (line 11). Reviewing (line 23), the browser selects, reads, integrates and evaluates information. The browse goal is modified (line 14) to continue review during the next browsing cycle in anticipation of further interest; for example, to learn more about a specific topic or to investigate a referential link. This browsing cycle continues as the browser reviews information and modifies goals.

3) Search-Oriented Browse: Even in a browsing setting, search-like behavior cropped up occasionally. The search-oriented browser plans search goals, scans to find relevant information, and reviews this information to integrate it. In search-oriented browse, the browser cycles through decisions that involve the methods *modify_search_goal, scan,* and *review.* Search-oriented browse is based on search goals; with goal modification based on the search goal operators stack goal and pop goal instead of the browse goal operator replace goal. He modifies search goals only as required to accomplish a fixed task. Sixteen of the subjects used search-browse, from one to five times each. Search-oriented browse is similar to the non-Boolean, goal-directed searches in literature review.

In terms of the BPM in Fig. 4, the search-oriented browser begins with a search goal to integrate information; to accomplish this goal, the browser stacks a search goal to find relevant information. The browser must first operate on the last goal stacked, to find information. Beginning with a search goal, the browser adopts the heuristic of relevance for the control strategy (line 5). The browser enters *pursue_goal* (line *7*)with the intention to find information (line 8), and selects *scan* (line 10). Scanning (line 19), the browser selects, reads, and evaluates for a match between information and goals. When the browser finds what appears to be relevant information, the browser enters *modify_goal* (line 12) to modify search goals

(line 14). Modifying search goals (line 28), the goal to find information is popped, having been satisfied. The browser then operates on the remaining goal to integrate information. The browser therefore reenters *pursue_goal* (line *7),* and selects *review* (line 11) in order to integrate information from the conceptual model with the browser's mental model. Reviewing (line 23), the browser selects, reads, integrates, and evaluates information. This browsing cycle continues (line 15) until the search goal is popped, having been satisfied or abandoned.

Some search-oriented browsers used only *scan* and not *review;* they stacked search goals just to find keywords for topics to browse. For example, in line 8 of Fig. 3, the subject searched the index using only *scan* to find information for further *scan* or *review.* With search-oriented browse, some amount of scanning is required to find information for further scanning or review. Other search-oriented browsers used both *scan* and *review;* they stacked search goals to both find and integrate information. For example, in the topic POW's, Subject #13 stacked a search goal to "see exactly where the troops were located since Nixon supposedly pulled everyone out." He decomposed it to find and review this information. He scanned and found the information ("that answered my question right there"). Then, with the search goal to integrate information, he reviewed the information ("they took off from aircraft carriers").

In summary, search-oriented browse can be defined as a process of search goal planning, scanning to find, and reviewing to integrate information relevant to a fixed task. Search-oriented browse is based on search goals. This differs significantly from the other two browsing strategies that are based on browse goals. In search-oriented browse, the goals are popped and stacked, decomposed instead of replaced. Their execution is planned instead of spontaneous. Their tasks are relatively fixed instead of freely variable. Information is evaluated for relevance instead of for interestingness. The distinction between strategies based on search goals and strategies based on browse goals is fundamental compared to the other differences between these strategies. Strategy based on search goals may be considered a possible variant in the browsing setting. Scan-browsing and review-browsing are the more conventional ways to browse. They involve scanning to find and reviewing to integrate interesting information in the presence of transient browse goals constrained by motivation. When this browsing involves a measurable amount of integration (of information) it is review-browsing; otherwise it is scan-browsing.

It should be understood that, in cognitive reality, any one of these browsing strategies might be interrupted or aborted in favor of a different task, different strategy, or other activity. The *Modify_goal* procedure is the common entry and exit point between these different tasks, strategies, or activities. The three browsing strategies identified are the components of complex browsing behavior.

VI. EXPERTS VERSUS NOVICES

This section addresses proposition #2, to examine differences in browsing behavior between subject-area experts

TABLE I
RESULTS OF DEPENDENT VARIABLES BY EXPERTISE

VAR NAME	EXPERT (10)		NOVICE (15)		t-TEST	
	MEAN	STD. DEV.	MEAN	STD. DEV.	t-VALUE	p
1 Topics Selected	7.40	1.26	9.87	4.34	1.74	0.10
2 Intertopic Links	0.40	0.52	3.47	3.44	2.18	0.02
3 Topics Aborted	0.60	1.07	5.00	3.80	N/A	N/A
Ratio: #3/#1	0.07	0.13	0.47	0.22	5.21	0.001
4 Special Features	0.90	1.52	1.13	1.55	0.31	None
5 Searches Implement	1.20	1.23	1.53	1.60	0.56	None
6 Searches Not Implem	0.70	1.51	0.87	1.81	0.24	None
7 Unknown Topics	0.00	0.00	1.07	1.22	2.74	0.02
8 Special Interests	N/A	N/A	N/A	N/A	N/A	N/A
9 Spec. Int. Topics	0.50	1.58	3.33	3.48	2.32	0.05
10 Expert Topics	6.20	1.14	4.87	3.70	N/A	N/A
Ratio: #10/#1	0.84	0.05	0.47	0.19	6.16	0.001
11 Integration Type:						
'A' : Add	0.9	0.32	0.73	0.46	1.00	None
'C': Change	0.50	0.53	0.33	0.49	0.91	None
'R': Reinforce	0.20	0.42	0.13	0.35	0.43	None
'T': Activate	0.20	0.42	0.07	0.26	0.98	None
12 Relating	1.10	1.37	0.27	1.03	1.74	0.10

and novices. Section **VI-A** discusses the use of the three browsing strategies identified in Section V — scan-browse, review-browse, and search-oriented browse — ordered by the number of subjects using each strategy. Section **VI-A** discusses the results of the detailed dependent variables.

A. *Overall Browsing Strategy*

1) Scan-Browse: Only two subjects scan-browsed: one was a novice and the other could not be classified. They usually scanned both index and topic nodes. They engaged in occasional review-browse, reviewing only ten of their combined 130 browsed topic nodes; however, scan-browsing was the overwhelming activity for these two subjects. Scan-browsing was an atypical browsing strategy in this experiment. One explanation for scan-browsing is that the subjects did not have persistent browse goals related to the subject area they were browsing; the transient browse goals they formulated were, therefore, so unmotivated that they did not inspire review. They were interested in very little of the material presented. Seldom did anything grab their attention.

Search-Oriented Browse: Sixteen subjects used search-oriented browse — six experts and ten novices — one to five times each. This is based on the number of persistent search goals observed in the protocol. Experts and novices did not differ significantly in their use of search-oriented browse (please see variable number 5 in). This strategy was used infrequently in the experiment. In the absence of specific direction, the subjects did not tend to generate specific goals of their own.

Review-Browse: Review-browsing was the predominant activity engaged in by 29 of the 31 subjects, experts and novices alike (of the remaining two, one was a novice and the other could not be classified). These subjects scanned index nodes

and reviewed topic nodes, modifying goals as desired, throughout their browsing session. Topics they browsed activated persistent browse goals and generated transient browse goals, inspiring review by both experts and novices.

Review-browse was clearly the predominant strategy used by both experts and novices — **29** of the **31** subjects. The remaining two subjects scan-browsed — one novice and one unclassified. The use of search-oriented browse was thinly scattered among the subjects. Based on these results, there was no difference in the overall browsing strategy used by experts and novices. It is necessary to look for differences in expert/novice behavior within overall browsing strategy; to make a detailed behavioral comparison by analyzing the individual dependent browse variables.

B. Detailed Dependent Browse Variables Analysis

Dependent variables were examined for subjects that could be classified as subject-area experts or novices based on the pretest, as previously described in Section III. Of the **26** subjects so classified, 25 were review-browsers and one was a scan-browser. In order to eliminate the overall browsing strategy as a new dependent variable, results from the scan-browser were disqualified from comparison. The remaining subjects are review-browsers — 10 experts and 15 novices. Dependent variables were measured in the four areas of interest: 1) browse path, **2)** search behavior, **3)** topics selected, and 4) mental models. Results are listed in Table I. For each dependent variable, Table I lists the number of observations, the sample mean, and standard deviation for experts and novices. **A** Student's two-sided t test (**23** d.f.) was used to test for significant differences between the expert and novice means. The table lists the resulting t-value and the probability of error (p).

A. *Browse Path:* Browse path variables, #1–#4, were examined for significant differences in the physical characteristics of paths taken while browsing. The results indicate that novices tended to skip from topic to topic more than did experts (variable #1, number of topics selected); novices made greater use of referential links than did experts (variable #2, referential links); and that novices aborted topics more often than did experts (variable #3, topics aborted). Novices appear to be less able or less willing to browse topics to completion. They are more prone to distraction by unknown terms. Novices want to assemble the basic nodes and links of their disjointed mental model into some kind of order; a process like assembling a jigsaw puzzle that changes with the placement of every piece. They used links built into the conceptual model to become familiar with entities and their relationships.

Experts, on the other hand, are already familiar with many basic entities and relationships. They appear to be more interested in more detailed knowledge of specific entities, as evidenced by their tendency to ignore referential links and review topics to completion. Experts have a more complete mental model with which to associate these greater levels of detail.

Special features, variable #4, counts the use of the links that surround the text area of topic screens. These were only used 28 times in total, nine times by three experts and 19 times by seven novices. They were not particularly attractive to either experts or novices. The t-test indicates no significant difference between experts 'and novices' use of these features.

B. Search Behavior: Search Behavior variables #5 and #6, *searches implemented* and *searches not implemented,* were examined for significant differences in tendencies to search between experts and novices. Very few searches implemented and searches not implemented were conducted by either experts or novices. The t-test indicated no significant difference between experts and novices in either of these variables. They did not progress from general to specific goals during the session. Their goals remained general, and only occasionally suggested search goals. Approximately 44% of search goals formulated were actually operated upon. The presence of a convenient search facility might have decreased the number of searches not implemented by making search convenient. In our experiment, browsers were forced to engage in search-oriented browse to find information, a relatively manual and time-consuming procedure.

C. Types of Topics Selected: Types of Topics Selected variables, #7-#10, were examined for significant differences in the type of topics selected. Unknown topics, #7, counts the number of topics selected that the browser has never previously heard of before — selection based on negative recognition. Only 16 of the 222 topics selected were unknown topics. Novice #16 selected five; she was characterized as a "determined novice" because she kept trying to integrate completely unfamiliar topics, aborting 15 or her 19 selections mostly through the use of referential links. It is more difficult and less interesting to integrate information into a mental model without links to recognized information; and the consequent effort requires more than a casual motivation. Interestingly, not one expert selected

an unknown topic, while 11 novices did. The t-test indicates with $p = 0.02$ that experts selected fewer unknown topics than did novices. One explanation is that this is consistent with the experts tendency, previously discussed, to be more interested in additional detail on known topics; compared to the novices interest in understanding basic relationships between topics. Another explanation is that experts are already familiar with so many topics that unknown topics are very exceptional; while for novices, unknown topics are the rule.

Special interests, #8 and special interest topics, #9, were examined to determine tendencies toward topics of this type. Special interests *identifies* subject areas of special interest. Special interest topics *counts* the number of topics selected that were related to the special interests. Novices selected 50 special interest topics, in contrast to only five selected by experts. A t-test indicates with $p = 0.05$ that novices browsed more special interest topics than did experts. Only one expert had a special interest that could be identified. His was Navy; he had taken a course in military history while in the Navy ROTC. Three of his six topics selected were related to the Navy — two by verbal association with the ROTC and one by reference to a "tribute to the people that were there," military and Navy personnel. In contrast, 11 of the 15 novices exhibited trends toward a special interest of some kind: people, the French influence on Vietnam, movies about Vietnam, politics, history, covert military activities, U.S. presidents, historic people, and hardships of war.

Novice #1, already discussed, was an excellent example of a browser with a special interest. Novice #27 is another. Her special interest was *hardships of war.* She worked toward this special interest in her first four topics, once saying "I want to get to the down and dirty stuff . . . I wish they'd show causes of death." She finally struck "pay dirt" in the topic WOMEN, which graphically portrayed the role of women treating the wounded and dying. About 30% of the topics selected by novices were related to special interests of browsers. The identification, naming, and counting of the number of topics related to special interest was subjective; but trends of this kind did seem pronounced enough to suggest that motivation often effected the overall direction of topic selection and integration in a persistent way. **An** explanation for the fact that novices had more identifiable special interests is that experts were browsing in an area of their special interest already — topics involving expert knowledge of The Vietnam War; while novices tried to find related areas of special interest to browse — topics of common sense knowledge or personal experience.

Expert Topics, #10, was examined to determine tendencies to select topics related to expert knowledge and reasoning as opposed to topics of common sense knowledge and reasoning. The t-test indicates with $p = 0.001$ confidence that experts browsed a higher percentage of topics related to expert knowledge and reasoning. As with special interests, the evaluation of expert topics was subjective; but trends in this direction seem pronounced enough to suggest that experts are more attracted to expert topics than are novices.

The results of the *Types of Topics Selected* variables suggest that experts and novices differ in the type of topics they are attracted to. Experts are attracted to topics involving expert

knowledge and reasoning. Novices are attracted to topics of special interest and common sense knowledge. Recognition is a common factor in almost all topic selections of any type for both experts and novices.

4) Mental Model: Mental Model variables #11 and #12, Integration Type and Relating, were examined for significant differences in the way experts and novices use and relate mental models. Integration Type, #11, is a subjective evaluation of the use to which integrated information is put: add to *(A)*, change (C), reinforce *(R)*, and activate *(T)*. As previously discussed (Section IV-B-3) *integrate),* no significant difference in integration type was observed between experts and novices. For example, experts did not try to reinforce their opinions more than novices; they were usually open-minded to new information as were most novices. This suggests simply that the basic use to which people put information is not dependent on subject-area expertise. This variable serves primarily to identify general uses of information, to indicate that integration of some kind was taking place, and to assist in modeling the integration process.

Relating, #12, was examined to determine whether experts relate differently to the conceptual model than do novices. *Relating* counts the number of times that browsers compare or contrast themselves and their mental models with the author of the Vietnam hypertext stack and his conceptual model. Six experts and one novice related to the mental and conceptual models as models. Examples of comments of this type are:

1) It seems to come across as pro-Westmoreland.
2) I want to see how the author felt about the subject.
3) I want to find out what they think about him.
4) Who prepared the text? It's one-sided.

The first three comments above were typical of the experts. The first represents an evaluation of the topic as "pro-Westmoreland." It is an evaluation that implies expert reasoning based on other expert knowledge of the topic; and that relates to "it," the conceptual model and its author, as a whole. The second comment relates directly to "how the author felt about the subject"; to the conceptual model and its author as a whole. The fourth comment that of a subject with a special interest in politics who believed he was an expert. He was primarily interested in reinforcing his opinions (integration type), and was prepared to do so by relating to the conceptual model and its author as biased. This shows that it is not necessary to be an expert to relate in this way.

Relating is characteristic of expertise or perceived expertise. It involves seeing the conceptual model as more than the sum of its information. The browser must have a fairly well developed model of a subject area in order to compare mental and conceptual models. Novices do not usually relate to their mental and conceptual models at this level of abstraction. They are still forming the nodes and links within their models, using the conceptual model for guidance.

Addressing the second research question, these results suggested that within the overall review-browsing strategy, the following differences between subject-area experts and novices were observed:

1) Experts tended to browse fewer topics but in more

depth than did novices. They tended to concentrate on obtaining more detail within known topics. They were not interested in changing topics until they completed review of the current topic.
2) Novices tended to rely more than experts on referential links. They tended to concentrate on understanding the basic structure of knowledge, and often switched topics to do so.
3) Experts tended to browse topics based on expert knowledge and reasoning. They were more interested in pursuing expertise, and saw the hyperstack as an opportunity to do so.
4) Novices tended to browse topics based on special interests and common sense knowledge. Without a strong motivation or persistent browse goal in expert areas, they turned to topics of general or special interest.
5) Experts tended more than novices to relate to the conceptual model at an abstracted level. They often compared it to their mental model and evaluated it for bias. They had, or believed they had, a sufficiently developed mental model to make such comparisons.

Results also suggested the following similarities:

1) The special features were not particularly useful to either experts or novices.
2) There was no overall tendency toward search-like behavior, or toward the formulation of specific goals. Browsers tended to remain with browse goals, occasionally formulating search goals to answer questions arising during browse.
3) Recognition was an important factor in topics selected. Subjects tended to browse topics with which they were already familiar.
4) Integration type (add, change, reinforce, activate) was not related to subject-area expertise.

VII. RECOMMENDATIONS TO HYPERTEXT DEVELOPERS

As a result of these findings, it is possible to make some recommendations to hypertext developers of similar systems. To do so, we must first describe "similar systems" at a comparable level of generality, as follows. The Vietnam War is a Macintosh Hypercard electronic hypertext application. It is designed to encourage browsing in a subject area for educational purposes (e.g., not business-oriented). It is intended to attract and retain both novices and experts in the subject area. It offers a set of indexes of keyword links organized under general headings; and it offers referential links between topics. The recommendations given below are applicable to the development of similar systems.

1) In general, design hypertext to facilitate recognition. Recognition helps to attract and retain browsers. Browsers are usually interested in topics they recognize. Most browsers base selections on positive recognition.
2) To attract and retain novices, provide topics related to their everyday, common sense knowledge. Examples of such topics in our experiment were NIXON and DRUGS. Novices were also attracted to topics of their special interest, such as controversial, political, human

interest, and historic topics. Build in topics of this type to attract and retain novices likely to have relevant special interests. These topics held novices' interest when expert topics did not.

3) Develop links from common sense and special interest topics into expert topics. This is the direction that novices work in; they often approach expert topics indirectly. This provides a natural association between new and existing knowledge that encourages the browser into expert topics. For example, with The Vietnam War hyperstack, a novice on the subject of Vietnam may browse the familiar topic Nixon, and in this context link to learn about a battle that would otherwise be of no interest. Linking peripheral topics into expert topics maximizes the usefulness of the peripheral topics by encouraging exploration into expert topics without disorienting the browser—the browser can safely return at any time to the peripheral, known topic.

4) Avoid losing browsers' interest by giving excessive detail in a node. Browsers may not want as much detail as reference material normally provides for study. It is better to retain interest longer by providing successively more detailed levels of information in separate nodes.

5) Experts like to be oriented to the information in the conceptual model. They are interested in the overall approach to the subject; where is it "coming from," what is "it's angle," who wrote it and why. Novices need to be oriented to the basic structure of the conceptual model. To help orient users, it is useful to view hypertext indexes as semantic networks. Design them to inform the user about the organization and structure of information as well as to provide the location of information. Classify topics in various ways and group them accordingly. These indexes will then help users create their mental models. For the expert, they should be designed to represent the conceptual model at appropriate levels of abstraction and classification. For the novice, they should help to understand the fundamental relationships between nodes and links.

VIII. Conclusion

We conducted this experiment to identify and operationally define browsing strategies, to compare differences in browsing between subject-area experts and novices, and recommend design guidelines to hypertext developers.

To operationally define browsing strategies, we first used a GOMS model to identify four basic types of goals (G): persistent search goals, transient search goals, persistent browse goals, and transient browse goals. Browse goals have a general objective to accomplish an open task while search goals have a specific objective to accomplish a closed task [30]. Observable browse goals are transient while observable search goals are persistent. Transient browse goals are goals that exist only temporarily, in short-term memory; they are normally replaced by other goals before being satisfied or intentionally abandoned.

Hill-climbing was the control strategy observed for browsing. The heuristic for evaluating information pertinent to search goals was *relevance* to a fixed task. The heuristic for evaluating information pertinent to browse goals was *interestingness* to changing tasks.

Six different operators (O) were identified and used to describe browsing behavior. These are *select*, a motor operator to find information; *read*, a perceptual operator to input information to the browser's mental model; *integrate*, a cognitive operator to alter the content or structure of the browser's mental model (add, change, reinforce, or activate); *evaluate*, a cognitive operator to obtain a match between information constrained by availability and goals constrained by motivation; the search goal operators *stack goal* and *pop goal*; and the browse goal operators *add goal, remove goal*, and *replace goal*.

Combinations of operators were observed to comprise methods (*M*) that were activated to achieve goals under certain conditions by selection rules (*S*). Two methods that involved motor operators were *scan* and *review*. *Scan* was used to *find* information; *review* was used to *integrate* information into the user's mental model. The browsing process can be described as the iterative application of two procedures: *pursue_goal*, choosing to scan or review information based on the intention to find or integrate it; and *modify_goals*, to set and change goals after evaluation.

The identification of these goals, operators, methods and selection rules allows us to view browsing as something more than just a random process. Transient browse goals are selected and pursued based on heuristics of interestingness, constrained by the relative influence of persistent browse goals (motivation) and the availability of information.

Three browsing strategies were operationally defined using the GOMS model. We call these *search-oriented browse, review-browse,* and *scan-browse.* Search-oriented browse is defined as the process of scanning to find and reviewing to integrate information relevant to a fixed task. Review-browse is defined as the process of scanning to find and reviewing to integrate interesting information in the presence of transient browse goals that are constrained by motivation. Scan-browse is defined as the process of scanning to find (and not reviewing to integrate) interesting information in the presence of transient browse goals that are constrained by motivation.

An analysis of the browsing strategies used by experts and novices revealed that both experts and novices engaged primarily in review-browsing (29 of the 31 subjects). In other words, we found no significant difference in the overall browsing strategies of experts and novices; however, a detailed examination of dependent variables revealed several differences in browsing tactics within these strategies. Novices used referential links more than experts; they needed to understand the basic structure of knowledge. Experts browsed fewer subjects, and in more depth than did novices; their familiarity with the basic nodes and links allowed them to concentrate on more detailed knowledge of specific nodes. Novices tended to browse topics based on special interests and common sense; they approached expert topics indirectly, through known topics. Experts tended to browse topics based

on expert knowledge and reasoning. Experts tended more than novices to evaluate the conceptual model for bias, as a whole; novices usually did not relate to the conceptual model at this level of abstraction.

Some similarities between the experts and novices were also observed. Recognition was an important factor in selecting topics for both experts and novices; very few selections were based on negative recognition. There was no overall tendency to move from general browse goals toward specific search goals. The use of information to effect the content or structure of the browser's mental model (add, change, reinforce or activate nodes and links) was not related to subject-area expertise.

Finally, we were able to make some recommendations to hypertext developers of similar browsing systems. To attract and retain novices, they should facilitate recognition by providing topics related to common sense and browsers special interests — topics peripheral to expert topics in the subject area. To encourage novice browsers into deeper knowledge, they should design links from these peripheral topics into expert topics. To maintain interest, they should provide successively more detailed levels of information in separate nodes. To help users relate to the organization and structure of the conceptual model, they should design the hypertext index as a semantic network that can serve as a foundation for users' mental models.

ACKNOWLEDGMENT

REFERENCES

[1] J. R. Anderson, *Cognitive Psychology and Its Implications.* New York: W. H. Freeman, 1985.
[2] M. J. Bates, "The design of browsing and berrypicking techniques for the on-line search interface," *Online Review,* vol. 13, no. 5, pp. 407–431, 1989.
[3] M. J. Bates, "Where should the person stop and the information search interface start?" *Information Processing and Management,* vol. 26, no. 5, pp. 575–591, 1990.
[4] C. L. Borgman, "Individual differences and information retrieval systems," *Information Processing and Management,* vol. 25, no. 3, pp. 237–251, 1989.
[5] C. L. Borgman, "The user's mental model of an information retrieval system: an experiment on a prototype on-line catalog," *Int. J. Man-Machine Studies,* vol. 24, pp. 47–64, 1985.
[6] J. M. Bouwman, "Human diagnostic reasoning by computers: **An** illustration from financial analysis," *Management Sci.,* vol. 29, pp. 653–672, 1983.
[7] S. K. Card, T. P. Moran, and A. Newell, *The Psychology of Human-Computer Interaction.* Hillsdate, NJ: Lawrence Erlbaum, 1983.
[8] A. Carlson, and S. Ram, "HyperIntelligence: The next frontier," *Commun. ACM,* vol. 33, no. 3, Mar. 1990.
[9] E. Carmel, W. K. McHenry, and Y. Cohen, "Building large dynamic hypertexts — How do we link intelligently," *J. Management Inform. Syst.,* vol. 6, no. 2, Fall, 1989.
[10] E. Charniak and D. McDermott, *Introduction to Artificial Intelligence.* Reading, MA: Addison-Wesley, 1985.
[11] H. Chen and V. Dhar, "Online query refinement on information retrieval systems: A process model of searcher/system interactions," in *Proc. 13th Int. Conf Res. Development in Inform. Retrieval,* (SIGIR-90), Brussels, Belgium, 1990.
[12] J. Conklin, "Hypertext: An introduction and survey," *IEEE Comput.,* pp. 17–41, Sept. 1987.
[13] J. F. Cove and B. C. Walsh, "Online text retrieval via browsing," *Inform. Processing and Management,* vol. 24, no. 1, pp. 31–37, 1988.
[14] E. Davis, *Representations of Commonsense Knowledge.* Palo Alto, CA: Morgan Kaufmann, 1990.
[15] S. J. DeRose, "Expanding the notion of links," *Hypertext'89 Proceedings.* Pittsburgh, PA: ACM Press, Nov. 1989, pp. 249–257.
[16] D. E. Egan, J. R. Remde, L. M. Gomez, T. K. Landauer, J. Eberhardt, and C. C. Lochbaum, "Formative design-evaluation of superbook," *ACM Trans. Inform. Syst.,* vol. 7, no. 1, pp. 30–57, Jan. 1989.
[17] K. A. Ericsson, and H. A. Simon, *Protocol Analysis: Verbal Report as Data.* Cambridge, MA: MIT Press, 1984.
[18] C. L. Foss, "Tools for reading and browsing hypertext," *Inform. Processing and Management,* vol. 25, no. 4, pp. 407–418, 1989.
[19] R. Gong, and J. Elkerton, "Designing minimal documentation using a GOMS model: A usability evaluation of an engineering approach," in *Proc. CHI'90: Empowering People,* 1990, pp. 99–106.
[20] F. G. Halasz, "Reflections on notecards: Seven Issues for the next generation of hypermedia systems.," *Commun. ACM,* vol. 31, no. 7, July, 1988.
[21] D. Hix and R. Schulman, "Human-computer interface development tools: A methodology for their evaluation," *Commun. ACM,* vol. 34, no. 3, pp. 75–87, Mar. 1991.
[22] B. R. Huguenard, M. J. Prietula, and J. F. Lerch, "Performance behavior: a study in the fragility of expertise," *Proc. Tenth Int. Conf Inform. Syst.* (ICIS-89), 1989, pp. 101–117.
[23] P. Jackson, *Introduction to Expert Systems,* second ed. New York: Addison-Wesley, pp. 19–21, 1990.
[24] B. E. John, "Extensions of GOMS analysis to expert performance requiring perception of dynamic visual and auditory information," *Proc. CHI'90: Empowering People,* 1990, pp. 107–115.
[25] B. E. John and A. Newell, "Cumulating the science of HCI: From S-R compatibility to transcription typing," in *Proc. CHI'89: Human Factors in Computing Systems,* 1989, pp. 109–114.
[26] D. B. Lenat, "AM: Discovery in mathematics as heuristic search," *Knowledge-Based Systems in Artificial Intelligence.* New York: McGraw-Hill, 1982.
[27] F. J. Lerch, M. M. Mantei, and J. R. Olson, "Skilled financial planning: the cost of translating ideas into action," in *Proc. CHI'89: Human Factors in Computing Syst.,* 1989, pp. 121–126.
[28] P. Liebscher and G. Marchionini, "Browse and analytical search strategies in a full-text CD-ROM encyclopedia," *School Library Media Quart.,* pp. 223–233, Summer 1988.
[29] G. Marchionini, "An invitation to browse: designing full text systems for novice users," *Can. J. Inform. Sci.,* vol. 12, no. 3, pp. 69–79, 1987.
[30] G. Marchionini, "Information-seeking strategies of novices using a full-text electronic encyclopedia," *J. Amer. Soc. Inform. Sci.,* vol. 40, no. 1, 1989, pp. 54–66, 1989.
[31] G. Marchionini and B. Shneiderman, "Finding facts vs. browsing knowledge in hypertext systems," *IEEE Comput.,* pp. 70–79, Jan. 1988.
[32] M. E. Maron, S. Curry, and P. Thompson, "An inductive search system: theory, design, and implementation," *IEEE Trans. Syst., Man, Cybern.,,* vol. SMC-16, Jan./Feb., 1986.
[33] R. S. Michalski, "A theory and methodology of inductive learning," *Artificial Intell.,* vol. 20, no. 3, 1983.
[34] A. Newell and H. A. Simon, *Human Problem Solving.* Englewood Cliffs, NJ: Prentice-Hall, 1972.
[35] J. Nielsen, *Hypertext and Hypermedia.* New York: Academic, 1990.
[36] J. Nielsen, "The art of navigating through hypertext," *Commun. ACM,* vol. 33, no. 3., pp. 298–310, Mar. 1990.
[37] D. A. Norman, "Some observations on mental models," *Mental Models,* 1983.
[38] J. R. Olson and E. Nilson, "Analysis of the cognition involved in spreadsheet software interaction," *Human-Computer Interaction,* vol. 3, pp. 309–349, 1988.
[39] J. R. Olson and G. M. Olson, "The growth of cognitive modeling in human-computer interaction since GOMS," *Human-Computer Interaction,* vol. 5, pp. 221–265, 1990.
[40] D. Ostroff and B. Shneiderman, "Selection devices for user of an electronic encyclopedia: an empirical comparison of four possibilities," *Information Processing and Management,* vol. 24, no. 6, 1988.
[41] P. G. Polson and C. H. Lewis, "Theory-based design for easily learned interfaces," *Human-Computer Interaction,* vol. 5, pp. 191–220, 1990.
[42] M. R. Quillian, "Semantic memory," *Semantic Information Processing.* Cambridge, MA: MIT Press, 1968.
[43] G. B. Salomon, "Designing casual-use hypertext — The CHI 89 Information Booth," in *Proc. CHI 90: Computer-Human Interface,* p. 451, 1990.
[44] B. Shneiderman, *Designing the User Interface: Strategies for Effective Human-Computer Interaction.* New York: Addison-Wesley, 1987.
[45] D. Waterman and A. Newell, "Protocol analysis as a task for artificial intelligence," *Artificial Intell.,* vol. 2, pp. 285–318, 1971.
[46] P. H. Winston, *Artificial Intelligence,* second ed. Reading, MA: Addison-Wesley, 1984, p. 132.

Erran Carmel received the B.A. degree from the University of California at Berkeley and the M.B.A. degree from the University of California at Los Angeles. He received the Ph.D. degree in management information systems in 1991 from the University of Arizona, Tucson.

He has been involved in information systems programming, management, marketing and sales in a number of firms including Xerox and American Management Systems. He was co-founder of a firm specializing in the audio-text market. He is currently Assistant Professor in the Department of Management at The American University, Washington DC, where his research interests include textual systems, electronic meeting systems, and software start-ups.

Hsinchun Chen received the Ph D degree in information systems from the New York University, New York, in 1989.

He is Assistant Professor of Management Information Systems at the University of Arizona. His research interests include human-computer interactions, text-based information management and retrieval, knowledge acquisition and discovery, and neural network modeling and classification He has published articles in such journals and periodicals as IEEE TRANSACTIONS ON SYSTEMS, MAN, AND CYBERNETICS. International Journal of Man-Machine Studies, Information Processing and Management, and Advances in Computers.

Stephen F. Crawford holds the B S degree in business administration from the University of Phoenix, and the M.S degree in management information systems from the University of Arizona, Tucson

He was a computer systems analyst for several years at IBM Corporation. He is currently self-employed as a consultant at Crawford Computer Consulting, specializing in ddtabase design and business systems integration