Cooperative Structuring of Information: The Representation of Reasoning and Debate

By David G. Lowe
Computer Science Department
Stanford University, Stanford, CA 94305
July 1984

Abstract

Interactive computer networks create new opportunities for the cooperative structuring of information which would be impossible to implement within a paper-based medium. Methods are described for cooperatively indexing, evaluating and synthesizing information through well-specified interactions by many users with a common database. These methods are based on the use of a structured representation for reasoning and debate, in which conclusions are explicitly justified or negated by individual items of evidence. Through debates on the accuracy of information and on aspects of the structures themselves, a large number of users can cooperatively rank all available items of information in terms of significance and relevance to each topic. Individual users can then choose the depth to which they wish to examine these structures for the purposes at hand. The function of this debate is not to arrive at specific conclusions, but rather to collect and order the best available evidence on each topic. By representing the basic structure of each field of knowledge, the system would function at one level as an information retrieval system in which documents are indexed, evaluated and ranked in the context of each topic of inquiry. At a deeper level, the system would encode knowledge in the argument structures themselves. This use of an interactive system for structuring information offers many further opportunities for improving the accuracy, integration, and accessibility of information.

This paper has since been published in the *International Journal of Man-Machine Studies,* **23** (1985), pp. 97-111.

Introduction

The availability of interactive computer networks creates many new opportunities for improving communication between groups of people who are working on a common problem. One of the largest and most important examples of this type of cooperative activity is the work of the worldwide scientific community. There are currently more than 60,000 scientific and technical journals being published regularly, as well as numerous conference proceedings, books, and reports. Frequently cited problems with these current methods of distribution include literature scatter, publication delays, rising costs, and inaccessibility. It would probably even now be economically advantageous to replace this paper-based distribution of scientific literature with electronic distribution through computer networks (Lancaster, 1978; Senders, 1976), and the economic advantage will continue to improve with rapidly falling computer and telecommunications costs. More significantly, electronic distribution could also provide important advantages in speed, flexibility, and retrieval capabilities (Shackel, 1982; Turoff, 1982). In particular, this paper will present new methods for structuring and retrieving information in interactive computer networks which could not be implemented within a paper-based medium.

Perhaps the most serious problem afflicting the current information distribution system is the problem of information overload. Except for narrow specialties, within which a researcher can continuously monitor all the relevant journals and conferences, the sheer size of the literature makes it very difficult to locate the most useful information on a given topic. There have been many calls and proposals in the past for the creation of encyclopedia systems which would synthesize information from many sources (Wells, 1938; Bush, 1945; Kochen, 1972; Pager, 1972; Soergel, 1977), but unfortunately these remain to be implemented. Currently, there are a number of large bibliographic retrieval systems in operation which allow a user to search for references with combinations of subject keywords (McCarn, 1980; Salton & McGill, 1983), but there are inadequacies in the use of keywords which make these systems difficult or unreliable to use. For example, most English words have ambiguous or multiple meanings and therefore do not precisely specify a subject topic. More importantly, these systems make no attempt to evaluate documents on the basis of accuracy, clarity, or other subjective criteria, although these criteria are of major importance in selecting which documents to use. The use of keywords to index documents is historically an outgrowth of the use of subject headings in library card catalogs, but there have been a number of more recent attempts to design new indexing facilities which make greater use of the computer's capabilities. In particular, there has been considerable interest in generalized methods for linking and referring to sections of text (Engelbart, 1973, 1984; Nelson, 1983; Trigg, 1983). These methods have proved to be very useful for facilitating user interaction and for creating networks of associations between related documents. However, the basis for these linked-text systems is still the individually-authored document. In this paper we will develop methods for synthesizing and structuring the content of many contributions as a way of directly attacking the problem of information overload.

Computers offer capabilities which would be almost impossible to provide through the traditional use of the printing press and the distribution of paper documents. They allow large numbers of people to interactively examine and modify a common body of information, which could only be carried out in a paper-based environment by continuously mailing updated versions following each modification. Computers also allow information to be structured far more flexibly than is possible within the linear ordering imposed by paper. These capabilities will be exploited in this paper to design a very different form of communication than those which are currently available. Rather than creating independent, individually authored documents which are then indexed and filed by editors and librarians, this new information resource would provide tools for users to interactively integrate their contributions and opinions into a single unified structure.

One major goal of this medium is to allow many users to interactively create a structured description of each field of study, within which documents can be referenced and evaluated according to the role which they play in that field. In this way the system would act as an up-to-date, extremely detailed textbook or survey article, evaluating and ranking documents according to their relationship to each topic within a field of study. However, another more radical goal of this medium is to combine the *content* of many contributions on any given topic into a single structure. Each item of information can be broken down to the level of individual sentences or concepts, and a well-specified set of user interactions with the system selects the best ordering and interrelationships between these concepts. By removing the redundancy between many individual contributions and representing each concept only once, the sheer quantity of information can be greatly reduced. By allowing many researchers to examine and suggest modifications to each structure, the accuracy, currency, and clarity of each presentation is likely to be much better than is possible with documents written by a single author.

The key development which allows many individuals to combine their thoughts and opinions on a topic is a representation for debate and the multiple viewpoints which can arise on any issue. The representation for reasoning and debate which will be described below requires each person to explicitly indicate their reasons for a given opinion, so that argument over a conclusion is transferred as much as possible to argument over the various sources of evidence for or against the conclusion. The purpose of debate is not to choose one answer to the exclusion of others, but rather to collect and order the presentation of evidence and to concisely summarize the range of opinion. Although a voting procedure is used to select the best candidates for initial presentation on any given topic, all contributions are retained and can be accessed if a topic is examined in sufficient depth. The explicit representation of debate allows many important but subjective matters—such as the significance of a topic or contribution—to be addressed, whereas it might not be politically acceptable to make these judgments in an information system which was created by a few individuals.

{A}	Overview: Information retrieval methods for access to document collections	
{B}	Keyword-based methods	[3]
{C}	Use of Boolean combinations of keywords for retrieval specification	[3]
{D}	Evaluation of keyword-based retrieval systems currently in operation	[2]
{E}	Keyword-based systems require a controlled vocabulary for accurate recall	[1]
{F}	Automatic generation of subject keywords from documents	[1]
{G}	Methods based on structural descriptions of document contents	[2]
{H}	Representation of reasoning and debate	[6]
{I}	Hierarchical representations of subject areas	[3]
{J}	Methods based on natural language understanding	[1]
{K}	Current capabilities of natural language understanding systems	[5]
{L}	Research on the use of natural language understanding for information retrieval	[2]
{M}	Predictions of future natural language capabilities of computers	[2]
{N}	More...	[1]

{O} Alternatives; {P} Search index; {Q} Back up; {R} Modify

Figure 1: The above display is shown to the user after a request for information on a specific topic such as "information retrieval." It presents an overview of the topic in familiar outline form. By typing the letters shown in braces (e.g., {B}), a user can examine any part of the structure in more detail or suggest modifications. For example, typing the letter "E" results in the display shown in Figure 2.

Examples of the system in operation

A version of this proposed information system has been implemented as a computer program and used for a number of experiments. The system has been named SYNVIEW to indicate its goal of combining multiple viewpoints into a single structure. Several examples of its use will now be presented to give the reader an intuitive feel for the representation before we embark on more theoretical issues.

Figure 1 gives an example of the first display from a computer terminal when a user asks for information on "information retrieval" (we will describe later how this request is made). This display presents an overview of the topic in familiar outline form—the sub-topics and sub-subtopics correspond roughly to what might be chapter and section headings in a textbook on the lead topic. However, the topics are ranked strictly in order of decreasing "importance" (in the sense of which topics are the most important to know for a general understanding of the lead topic) rather than by any of the other criteria that are often used in writing. The number

{A}	Keyword-based information retrieval systems require a controlled vocabulary for accurate retrieval	[5]
{B}	Most English words have imprecise or multiple meanings	[6] ↑[4]
{C}	A controlled vocabulary is needed for precoordination of index terms	[6] ↑[4]
{D}	There are usually many approximately synonymous words for any given topic	[8] ↑[1]
{E}	However: Within small specialized technical domains the natural vocabulary may have adequate precision	[2] ↑[-1]
{F}	All large commercial bibliographic retrieval systems have chosen to use a controlled vocabulary	[4] ↑[1]
{G}	Can accurate retrieval be achieved through statistical operations on ambiguous keywords?	[-1] ↑[4]
{H}	More...	[0]

{I} Alternatives; {J} Search index; {K} Back up; {L} Modify

Figure 2: This display is an example of the top-level structure of debate. The line labelled {A} is the assertion or question for which the following lines are items of evidence in decreasing order of importance. Each item of evidence can be individually selected to examine its own support or to debate the relationship between the evidence and the assertion. For example, selecting {B} produces the display shown in Figure 5. Selecting the top line (i.e., {A}) produces introductory information on the topic as shown in Figure 6.

in square brackets to the right of each subtopic gives its rated importance with respect to the next higher level. The letters in braces at the beginning of each line (e.g., {A}) are menu selection terms—by typing a given letter the user can descend in the hierarchy of topics to examine any subtopic in more detail. SYNVIEW does not display all topics at one level before displaying any topics at the next level; rather, the display is balanced so that the cutoff in importance is at a constant value with respect to the head topic. In order to see more of the top-level topics, the user can select the line labeled "More."

The creation and modification of these information structures is based upon the input of many individuals. Given the display shown in Figure 1, a user can suggest modifications or additions along any of a number of dimensions. At the simplest level, a user can give his or her opinion on the importance of a subtopic and the vote will be averaged in with the others in determining the ordering of the subtopics. If the user disagrees with the wording used for some topic, he or she can suggest a different wording to be used. The choice between alternative wordings then becomes a topic for debate (as described below), with different users entering and voting on reasons as to why one wording is superior to another. Any user can also add new subtopics below any topic, although it could be ranked far down on the list if others judge it to be unimportant or irrelevant. When suggesting a new wording for a topic, it is possible to create an entirely new set of subtopics below it and thereby completely redesign the organization of the presentation. The choice between these alternative organizations also becomes a topic for debate.

The "overview" structure shown in Figure 1 is useful for listing all the relevant subtopics for some field of study. However, a type of structure which is much more central to this medium of communication is the representation of debate. Most of the items shown in Figure 1 are noun phrases naming general topics of discussion. However, the line labelled {E} is phrased as a sentence and statement of fact, signaling the start of a structured debate. In such a debate, all subtopics are further statements giving specific items of evidence in support of or against the topic of the debate. Figure 2 shows the display that is presented when the user selects the item "E." Each item of evidence can be questioned in two different ways: the evidence itself may be judged true or false to varying degrees, and the implication of the original assertion may or may not follow from the truth of the evidence (the two numbers to the right of each line of evidence refer respectively to these judgments). Of course, each item of evidence is itself an assertion which can be examined in the same way as the original one. In addition to representing much of the material of any scientific discipline, these debate structures are used within SYNVIEW to resolve various areas of disagreement during the creation of structures. The next section of this paper will examine the representation and use of debate in greater detail.

In the above examples, all the information was presented in outline form with the information broken up into short phrases or sentences. This representation has some advantages in terms of compactness and the ability to skip to any part of a structure without reading the previous context, but it lacks some of the flow, legibility and tutorial nature of normal writing. To a certain extent, the use of outline form is a price that must be paid for the capability of having many people interactively and incrementally contribute to a single structure—the structure must be broken down into individual concepts so that every aspect can be questioned and debated. In addition, the fact that phrases may be reordered or have alternative versions substituted for them means that many of the mechanisms of normal writing which rely on ordering must be avoided. However, there are also mechanisms for integrating sections of text with the outline form that can be used for those situations in which normal text is preferable. It is possible to refer to on-line or off-line textual material as part of the presentation on any topic, and debates can be used to rank the documents and discuss their merits. In particular, the use of outline form seems better for structuring material for presentation to users who are familiar with a subject than for presenting tutorial material. Therefore, there is a uniform interface to introductory tutorials on each topic, mostly in document form, which will be described later. In fact, many people could contribute to a single textual document by cooperatively creating it in outline form, having a single writer produce a textual version, and then debating possible modifications to the document.

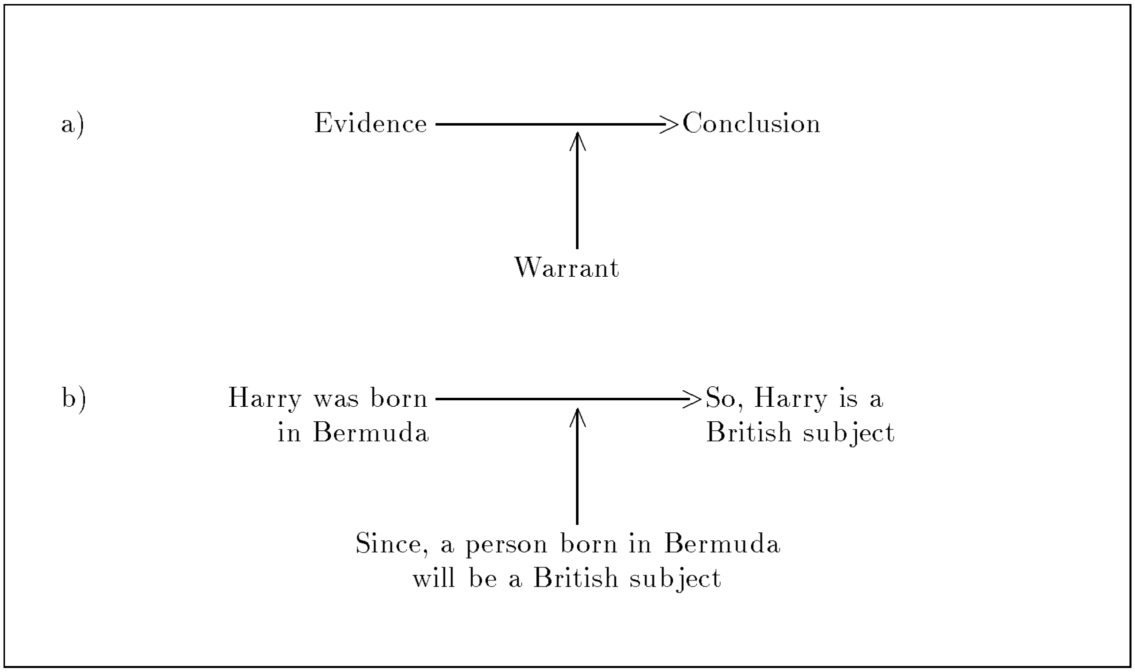

Figure 3: This example from Toulmin shows the most basic form of argument, in which evidence is presented for a conclusion based on an (often unstated) warrant. This is similar to the syllogistic form of argument typically studied by logicians. The diagram in (a) is the general form, and (b) shows a specific example.

The representation of debate

The explicit representation of the evidence and reasoning involved in reaching conclusions is the most important requirement for the development of this system. Without a representation of the reasons for reaching a conclusion, many aspects of the structure would amount to little more than opinion polls on the accuracy of some statement. By representing the evidence and forcing users to point to which evidence they are using for their conclusions—and their reasons for discounting contrary evidence—the user will be able to compare his or her own judgments on individual items of evidence to those of the other users of the system. There is also some reason to believe that the full presentation and consideration of evidence relating to any given conclusion will lead to less conflict of opinion than is common with current information distribution methods. It should be emphasized again that the purpose of representing debate is not to reach absolute conclusions, but rather to collect and order the best available sources of evidence for each significant issue.

Fortunately, there has already been an extensive amount of research into the representation of human reasoning. In particular, philosophers have been debating the structure of reasoning and inference since the time of Aristotle. Unfortunately, the model of "deductive inference" studied by most philosophers treats highly idealized cases (e.g., "all men are mortal") and fails to capture the form of most human debates, which are not subject to absolute proof. However, there has been some more practical work, such as the work of the philosopher Stephen Toulmin on the

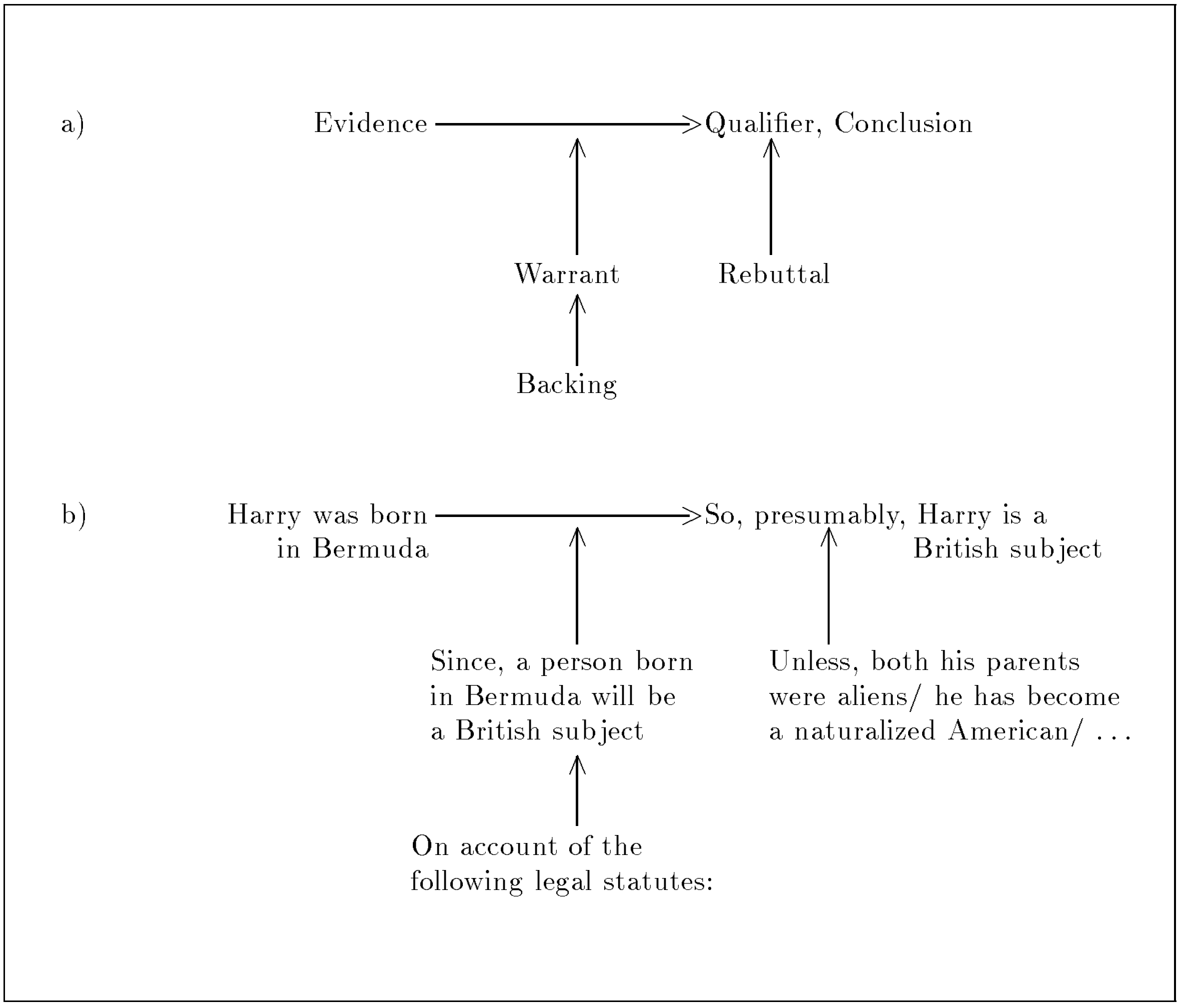

Figure 4: Human reasoning is seldom based on absolute proof, and there are typically an indefinite number exceptions to any rule. Toulmin introduced the concepts of a qualifier (which indicates the universality of an argument), a rebuttal (which gives conditions of exception to the warrant), and backing (which gives evidence for the warrant). Once again, (a) is the general form and (b) is an example.

layout of arguments. In his widely acknowledged book, *The Uses of Argument,* (Toulmin, 1958) questions the usefulness of traditional work on logic and deductive inference, and proposes a practical form of structuring argument as shown in Figures 3 and 4. Figure 3 shows an argument resembling the traditional syllogism introduced by Aristotle and used in most work on deductive inference. However, very few human arguments can be fully cast within this framework due to the indefinite number of exceptions and counterclaims which can typically be brought to bear on any argument. For this reason Toulmin introduced a number of other components to the layout of an argument, including the qualifier, rebuttal and backing, as shown in Figure 4.

The representation of arguments in SYNVIEW follows Toulmin's general structure and terminology, but also makes a number of simplifications. We saw in Figure

2 how an argument is represented at the first level by a ranked list of items of evidence for and against a particular conclusion. However, that list contained no explicit consideration of the warrant linking each item of evidence to the conclusion. When one of the items of evidence is selected from a display such as that in Figure 2, a new type of display is created, as shown in Figure 5, in which the original conclusion is introduced with the word “Context” and the warrant is shown explicitly and justified with its own evidence. In this way, the warrant becomes just another conclusion for which items of evidence can be presented. These items of evidence for and against the warrant combine Toulmin’s backing and rebuttal categories, and the use of numbers for evaluation correspond to his qualifier. The example in Figure 5 has been made somewhat more complicated than the typical case for the purposes of demonstration. It is usually more straightforward to reason with steps which place most of the debate under the evidence rather than the warrant, since it can be tedious to debate the strength of an implication rather than the truth of a more concrete assertion. Toulmin himself said that in normal debate evidence is usually appealed to explicitly, warrants implicitly.

The ranking of evidence

Given an understanding of warrants, it is possible to explain in more detail the function of the numbers in square brackets at the end of each line. These numbers are the median values of the many votes cast on the degree of truth or correctness of the statement (as implied by the evidence under it). The votes are given on a scale ranging from −10 (for false with no possible doubt), through 0 (for no idea whether true or false), to 10 (for true with no possible doubt). At the moment, the intermediate point of 5 has been pegged as the point at which it is estimated that there is only a 1% chance that the statement is false. In order to assure the maximum agreement and repeatability across different users, it would be necessary to develop a more detailed scale giving guidelines for assigning these numbers. Note that these numbers are also intended to capture any ambiguity and imprecision of statements. A median value is used rather than an average because when taking an average someone can always increase the influence of their vote by moving it towards one of the extremes—the median does not encourage overcompensation. Other facilities could be provided for presenting a more detailed histogram of votes should a user wish to examine them in more detail.

The first number following any assertion refers to the evaluation of that item. The second number following an item of evidence is preceded by an upward arrow (e.g., ↑[3]) and refers to the truth of the warrant linking that item of evidence to the stated conclusion. In other words, if the evidence was definitely true then this is the strength with which it would imply the truth of the conclusion. For an item of evidence to be significant, both its own strength and the strength of its warrant must be high. Therefore, items of evidence are automatically ranked in decreasing order of the minimum of these two values.

The values following items under an “overview” heading, as in Figure 1, are ranked on a different scale, since these items are not true or false but merely more or

{A}	Context: Keyword-based information retrieval systems require a controlled vocabulary for accurate retrieval	
{B}	Most English words have imprecise or multiple meanings	[6]
{C}	From 60 to 80% of English words have more than one currently used meaning	[6] ↑[7]
{D}	At least 10% of English word definitions vary substantially according to subject area, geographical region, or educational level	[4] ↑[3]
{E}	Experiments show large individual differences in the categorization of some objects according to common nouns	[6] ↑[1]
{F}	↑Warrant: The ambiguity of English words implies that the choice of keywords must be controlled	[4]
{G}	If a keyword has meanings other than the intended one, it could result in the incorrect retrieval of items specified by the other meanings	[7] ↑[5]
{H}	However: Technical words and compound words are less ambiguous than English words in general	[7] ↑[-2]
{I}	However: The retrieval system could engage in a dialogue with the user to select the appropriate meaning for an ambiguous keyword	[3] ↑[-2]

{J} Alternatives; {K} Search index; {L} Back up; {M} Modify

Figure 5: This display examines in greater detail the reasoning from an item of evidence to a conclusion. The conclusion is given with the heading "Context" in line {A}, the item of evidence is given in line {B}, and a paraphrase of the warrant linking the evidence to the conclusion is given in line {F}. The evidence and warrant are then justified by their own items in evidence in the remaining lines.

less significant with respect to each other. Therefore, each item is assigned a relative significance such that the significance of all items adds up to 10. The significance of an item can be thought of as roughly proportional to the amount of space that would be devoted to it in a short survey article on the overview topic. Note that each overview item is expressed as a noun phrase, whereas the items which are conclusions are expressed as complete sentences.

For many people, the voting on the truth of conclusions may seem to be rather unsatisfactory, since truth is not something which is based on popularity. However, the voting is not meant to establish some final truth, but rather to indicate the range of opinion and to rank items in the most useful order for the other people who will examine the structure. We are looking for order-of-magnitude estimates of an item's importance, and small deviations have little effect on the final presentation. Although it has not yet been implemented in the current system, there could be some kind of enforced consistency between votes on evidence and warrants and votes on conclusions. The system could also enforce the requirement that someone indicate their reasons for a particular judgment, so that others may question it and produce countering evidence.

Maintaining consistency during modification

The most common methods for making modifications are to use the voting procedure to change the rankings of items of evidence or to suggest new items to be added to a structure. We have also mentioned the possibility of suggesting a change of wording for some item, which results in a debate comparing the suggested wording with the current one. However, there are other occassions in which a number of changes must be made simultaneously if the overall presentation is to remain consistent. One example is a change in structural layout, in which a new category may be added and several current categories moved underneath it. A more extensive example is a change in terminology, in which the wordings of many items should be changed simultaneously in order to maintain consistent usage of the terminology. In each of these examples, people may have different opinions regarding the desirability of the change, but they may agree that either a complete set of changes should be made or none should be made at all.

Therefore, a facility is needed for grouping sets of changes together and debating the desirability of the set as a whole. In the current implementation, there are three types of changes that can be specified: (1) adding a new item under some heading, (2) moving a current item from being under one heading to another, or (3) replacing the wording of one item with another. A user can assemble a list containing any number of these changes and create a single debate regarding the desirability of making this set of changes. Note that these debates regarding changes in structure will actually result in the changes being made by the system if they receive positive votes and certain refereeing conditions are met as described below. On the other hand, the debate structure and a description of the inverse change are still retained, so we have maintained the rule that only the order of presentation is affected and all information input to the system is available at some level.

With large numbers of users contributing to the system, it is possible that these change structures will become numerous and complex. However, casual users need never look at any of these structures if they do not so desire. For the specialist in some field or for a person who questions the way in which a topic is presented, these debates regarding presentation and structure could be of substantial interest in themselves. All of the changes for each topic (accessed by selecting the "Alternatives" menu item) are organized in an overview structure of the type shown in Figure 1, so that changes can be ranked in order of importance and can even be structured into sets of related types of changes. These mechanisms are all necessary to achieve our objective of making explicit all of the normally implicit considerations that go into presenting information on a particular topic.

```
{A}  Introduction: Keyword-based information retrieval systems
        require a controlled vocabulary for accurate retrieval
{B}       Definitions
{C}             A controlled vocabulary is a list of allowable           [4]
                   words with a single definition specified for
                   each word
{D}             Retrieval is considered accurate if all relevant         [3]
                   items and only relevant items are recalled by
                   a knowledgeable request
{E}       Introductory references
{F}             [Controlled vocabularies for information retrieval       [6]
                   (350 words)] {G}
{H}             [Salton and McGill, 1983, Chapter 4 (4000 words)]        [1]
{I}             [Lancaster, 1976, Chapter 2 (2000 words)]                [1]
{J} Search index; {K} Back up; {L} Modify
```

Figure 6: This display provides tutorial material to introduce a user to an unfamiliar topic. Short definitions of new words or concepts are given, and a number of introductory references are suggested. The first reference is written expressly for use in this retrieval system, and is available online by selecting {G}. Of course, further information and discussion on any definition or reference can be obtained by selecting the line it is on.

Interface with tutorial materials and documents

A basic function of information retrieval is to not only answer specific questions but also to teach the user what he or she must know in order to understand material in some subject area. As mentioned earlier, the outline form of presentation seems less useful for tutorial purposes than for those who already have some understanding of the material. For this reason and for purposes of integrating the system with current methods of information distribution, a standard interface is provided between the outline structures and introductory material or other documents on any given topic. When the first line (the {A}) of a display such as those in Figures 1 or 2 is selected, the system has already displayed the subtopics for that line and therefore switches to an "Introduction" display as shown in Figure 6. This display contains definitions for any words or concepts not dealt with in higher-level displays and references to the most useful tutorial documents on the topic (evaluated according to their tutorial value). This section could be expanded in numerous ways: there could be interfaces to computer-aided instruction programs, lists of examples or problems to solve, names of experts in the field who would be willing to answer questions, access to multi-media courses on the topic, etc.

A related aspect of information retrieval is to vary the presentation of information according to various characteristics of the user, such as the user's prefered language, educational background, or areas of expertise. A simple example—but one of major importance for research communities—is to have the system present information in the user's language of choice. Translations between languages could be provided by many users and the choice of best translation would be a topic for debate. Another example would be to have a different presentation for someone with extensive knowledge of a field as compared with someone who knows little about it. Each topic could have several different presentations depending on the user's background, and the users would characterize themselves as having different degrees of knowledge in given areas. It should, of course, be easy to switch a display between different levels of presentation. This capability for varying presentations has not yet been incorporated into SYNVIEW, but it would not be difficult to do so.

Of course, references to individually-authored documents can appear in many other parts of the information structure than just the tutorial sections. The evidence for any conclusion can consist of experiments, statements, or eyewitness testimony reported in traditional documents. The documents and the claims they make would become further topics for overview and debate structures. In this way, documents can be integrated and indexed in a natural way. Ideally, the documents would be available on-line for immediate access. In addition to textual documents, there would be considerable value in similarly integrating graphics, tables, pictures, and multi-media material as the technology allows.

Index for initial access

The examples above illustrated how retrieval is accomplished by traversing the structures describing the relevant area of knowledge. As the user becomes familiar with the representation of the subject area, it should become continuously easier to find some desired item of information. However, this still leaves the question of how a search is initiated at the most relevant starting point in a potentially very large base of knowledge. There would probably be some use in having a hierarchical description of the entire knowledge base for the purpose of browsing and seeing what is available, but there must also be much more direct methods for accessing any particular structure.

The indexing method used in SYNVIEW uses keywords and modifying phrases as illustrated in Figure 7. The display is much like that typically used for the index at the back of a book, but with a few differences. One difference is that the modifying phrases under any indexing phrase are not in alphabetical order. In general, there is never a need to put things in alphabetical order in a computer, since the purpose of alphabetical order is to allow easy searching for a specific item, which is a task that the computer can perform directly. More importantly, the computer can keep track of how often the various modifying phrases are selected and rank them in decreasing order of frequency of selection. This minimizes the average distance that a person must search to find the phrase of interest, and allows the list to be of indefinite length without increasing the average search time.

As with all other aspects of the system, the index is constructed interactively by all the users of the system. Not only can users suggest indexing terms when they create a specific structure, but they can suggest new terms whenever they use the index and find that some index term they attempt is not present. If there is

	Information retrieval {A}
{B}	Computer access to documents {C}
{D}	Computer access to databases {E}
{F}	Traditional library methods for {G}
{H}	Future of {I}
{J}	Economic justification of {K}
{L}	Current systems for {M}
{N}	Keyword-based methods for {O}
{P}	Computer use of natural language for {Q}
{R}	History of {S}
{T}	More...

Figure 7: This section of the index is displayed when the user enters an index request for "information retrieval." The menu items on the left search the index to a greater depth, while those on the right move to the appropriate location in the information structure dealing with that topic. Items are ranked by decreasing frequency of access.

disagreement as to which structure is the most relevant for some index term or which wording should be used for a topic, these can become topics for debate.

Social considerations

Products of the current information system are each produced by at most a small number of authors. This paper has described alternative methods in which an indefinitely large number of contributors can create a coherent and concise result which attempts to take advantage of the best of the many contributions. The problem now arises as to whether we have gone too far to the opposite extreme—giving experts who have the deepest and most reliable knowledge of a topic no more influence than any other user of the system. This is not a problem in terms of who contributes new items of information—any low-quality contributions can be ranked far down in the order of presentation where they will have little influence—but it is definitely an important consideration in terms of who selects the rankings. However, there are a number of factors which make these problems less serious than they would be in other information systems. Rather than making arbitrary pronouncements on the correctness and significance of each statement, users are almost always constrained to indicating the strengths of implications from various items of evidence to conclusions. It is reasonable to expect that less expertise will be required and that there will be less disagreement in judging the strength of these implications than there would be for judging correctness in isolation. In addition, to the extent that the system is successful in bringing together all the relevant considerations and information needed to reach a conclusion, a much larger number of people will be able to become "expert" in a given topic than is currently the case.

However, in spite of these factors, it still remains the case that most people will probably want to see what "the experts" or some other chosen group think about some topic. Fortunately, computers allow us to satisfy everyone's opinions on this question by keeping track of many different voting groups for each topic and allowing users to choose and compare among them. There could be a general facility which allows users to view structures as created by various specified groupings of individuals or even by just a single person. Although this has not been implemented in the current system, the dynamic arrangement of structures does not appear to be computationally impractical. The objective should be to allow cooperative structures when desired rather than to force them upon everyone.

A related issue is whether the identities of users and their votes on specific topics should be publically accessible. This is the case with most current distribution methods, since almost all current academic writing is publically identified with a particular author. Having identities available allows people to examine and comment on important issues such as conflict of interest which may influence voting behavior. On the other hand, there are some opinions that people would avoid expressing if their names were publically attached to them. Once again, the best answer is probably to leave the choice of anonymity up to the individual users. Of course, the system itself must still maintain some sort of record of identities in order to prevent double voting. Given a record of this information, other capabilities of interest such as tracing the areas of agreement and disagreement leading to a conflict in opinion and examining correlations in opinions can probably be done without revealing the identities of individual contributors.

One case in which users definitely want their names to be public is when they are looking for credit for being the originator of some idea. In fact, a major function of current academic literature is to establish claims and credit for research results. The system can of course maintain a record of the original contributor of each idea, and this record can be debated by other users as to which contributions were the most important for the development of some new result. Various other mechanisms can be devised to acknowledge or reward users for contributing to the system, just as there is currently a reward structure for writing research papers or serving as a journal editor or reviewer.

There are difficulties in introducing new ideas to already-built structures which require a formalized peer-review process similar to the process a paper must go through before being accepted to a journal. The problem is that when someone introduces a new concept, his or her evaluation of its significance with respect to other ideas is not sufficient for ranking it among much more carefully evaluated ideas which already exist. Therefore, new ideas must first be subject to a peer review process in which a number of other (preferably randomized) users are asked to evaluate the rankings. This provides a more reliable statistical sampling of user opinion before elevating an item to undue significance. The more significant the initial ratings are, the larger the number of users that should be requested to evaluate the idea. The maintenance of lists of willing reviewers and the distribution of requests for evaluations could all be done automatically by the system. In any case, automatic notification of important new contributions and results in specified fields of study is likely to be a widely used facility. Old votes should expire as important new evidence or results are brought to bear in a field, and the original contributors should be notified of the need to update their input. The use of peer review and automatic notification can provide almost immediate feedback to contributors, and provide a public record of debate on the latest scientific issues. Currently, much of this debate is carried out in private or has a cycle time measured in years in published journals or books. The popularity of the few attempts at publically accessible review and debate (Harnad, 1982) indicates the value of this information.

One of the potentially most significant social impacts of a system like SYNVIEW is also one of the most difficult to measure. This is the potential for aiding in conflict resolution. Many strong conflicts in opinion can be traced to a reliance on different sources of information which are themselves written from the viewpoints of their readers. There is therefore some reason to believe that conflict would be lessened by exposing users to all available evidence and by making the user justify his or her evaluations of a conclusion in terms of each item of evidence. The use of a range of values for evaluation provides ample middle ground for consensus where evidence is lacking or contradictory. Only experience with a working system will tell how strong these effects are.

Technological factors

A major factor in the level of use of network-based interactive computer systems is cost. Current bibliographic retrieval systems often charge $30-70 per hour for use of the systems, and these costs are projected to double over the next few years even as computer hardware costs continue to fall. However, the major component of these prices is the cost of creating the databases—costs that would not exist for information entered free of charge by users in an academic environment. In addition, these systems use outdated hardware architectures which service all users from a single location on machines which are far more expensive per function than currently available technology.

A much cheaper implementation—which would also allow for indefinite expansion of the system—would be to have many individual nodes geographically distributed to be near the users. Each node would have a moderate computational and storage capacity and would store locally the most frequently accessed information. The less frequently accessed information would be distributed among different nodes and accessed automatically over telecommunication links between the nodes. Telecommunications costs between nodes should be well under a dollar per user per hour of use. If a node is locally available, telecommunications costs between users and nodes would be negligible. The most substantial part of the cost of each node would likely be the long-term online storage. It is difficult to estimate the total amount of information which must be stored in the entire system. However, it would not be prohibitive to store even the entire scientific literature in a widely used network, and given the removal of redundancy inherent in the SYNVIEW representation, storage requirements are likely to be far lower. The introduction of digital storage on optical disks will soon lower the cost of long-term online storage by a large factor.

Summary

There have, of course, been previous attempts to collect information from many contributors into a common resource (e.g., encyclopedias, collections of papers, peer commentary, electronic bulletin boards), but these have been limited in scope and usefulness by the amount of time that it takes readers to wade through the accumulation of material. In contrast, this paper has outlined methods for synthesis and selection which can operate on indefinitely large quantities of information and yet present them in manageable structures which can be examined to whatever depth the user desires. These methods operate by structuring and ranking information so that each structure starts with the most important and reliable items of information on that topic. Perhaps of most importance, these methods leave an explicit trail of the collective reasoning that went into arriving at each conclusion, so that any new user will be able to compare his or her use of the available evidence with that of those who have provided previous input to the system. The use of these facilities could result in a major improvement in the availability of accurate, clear, concise, and up-to-date information.

There is a preliminary implementation of SYNVIEW which produced the various examples shown in this paper. Personal experience by the author with the working system was a major factor in many design decisions, and no doubt more extensive use will suggest further modifications. The system has proved useful even for structuring the information and notes of a single user. Conclusions regarding its performance in large communities of users must wait until more extensive trials can be performed. Of course, use of an information system of this type need not be confined to the sciences or even to academic discussion. Examples of other applications include providing consumer information in which products or services are evaluated, predicting and evaluating the effects of legislative or other proposals before they are implemented, and distributing updates on current events as a component of the news media.

Possibly more significant than the actual system which has been described is the concept of having many people cooperatively build a common structure containing the best of their many contributions. If nothing else, I hope this paper has persuaded the reader of the importance and potential of this topic.

References

Bush, V. (1945). As we may think. *Atlantic Monthly,* **176,** 1, pp. 101-108. (Reprinted in Kochen, M. (ed.) *The Growth of Knowledge,* New York: John Wiley & Sons, 1967, pp. 23-35.)

Engelbart, Douglas C., R. W. Watson & J. C. Norton (1973). The augmented knowledge workshop. *Proceedings AFIPS National Computer Conference,* **42** (Arlington, Va.), pp. 9-21.

Engelbart, Douglas C. (1984). Collaboration support provisions in AUGMENT. *AFIPS 1984 Office Automation Conference,* Los Angeles, (February).

Harnad, Stevan, Ed. (1982). Peer commentary on peer review. *The Behavioral and Brain Sciences,* **5,** pp. 185-255.

Kochen, Manfred (1972). WISE: A world information synthesis and encyclopaedia. *Journal of Documentation,* **28,** 4, pp. 322-343 (December).

Lancaster, F.W. (1978) *Toward Paperless Information Systems.* New York: Academic Press.

McCarn, Davis B. (1980). MEDLINE: An introduction to on-line searching. *Journal of the American Society for Information Science,* pp. 181-192 (May).

Nelson, Ted H. (1983). *Literary Machines.* (Available from Ted Nelson, Box 128, Swarthmore, PA 19081).

Pager, David. (1972). A proposal for a computer-based interactive scientific community. *Communications of the ACM,* **15,** 2, pp. 71-75 (February).

Salton, Gerard & Michael J. McGill (1983). *Introduction to Modern Information Retrieval.* New York: McGraw-Hill.

Senders, John (1976). The scientific journal of the future. *The American Sociologist,* **11,** pp. 160-164 (August).

Shackel, B. (1982). Plans and initial progress with BLEND—an electronic network communication experiment. *International Journal of Man-Machine Studies,* **17,** pp. 225-233.

Soergel, Dagobert (1977). An automated encyclopedia—a solution of the information problem? *International Classification,* **4,** 1, pp. 4-10 & 2, pp. 81-87.

Toulmin, Stephen Edelston (1958). *The Uses of Argument.* Cambridge University Press.

Trigg, Randall H. (1983). *A Network-Based Approach to Text Handling for the Online Scientific Community.* Ph.D. Dissertation, Department of Computer Science, University of Maryland.

Turoff, Murray and Starr Roxanne Hiltz (1982). The electronic journal: A progress report. *Journal of the American Society for Information Science,* pp. 195-202 (July).

Wells, H. G. (1938). *World Brain,* Doubleday, Doran & Co., Inc., Garden City, N.Y. (Reprinted in Kochen, M. (ed.) *The Growth of Knowledge,* New York: John Wiley & Sons, 1967, pp. 11-22.)