A Prototype Electronic Encyclopedia

STEPHEN A. WEYER
Atari Sunnyvale Research Laboratory
and
ALAN H. BORNING
University of Washington

We describe a prototype electronic encyclopedia implemented on a powerful personal computer, in which user interface, media presentation, and knowledge representation techniques are applied to improving access to a knowledge resource. In itself, an electronic encyclopedia is an important information resource, but this work also illustrates the issues and approaches for many types of electronic information retrieval environments. In the prototype we make dynamic use of the structure and semantics of the text articles and index of an existing encyclopedia, while experimenting with other forms of representation, such as simulation and videodisc images. We present a long-term vision of an intelligent user-interface agent; summarize previous work related to futuristic encyclopedias, electronic books, decision support systems, and knowledge libraries; and outline current and potential research directions.

CR Categories and Subject Descriptors: H.3.4 [**Information Storage and Retrieval**]: Systems and Software

General Terms: Experimentation, Human Factors

Additional Key Words and Phrases: Electronic books, information browsing and viewing, interactive simulations, videodiscs

1. INTRODUCTION

Encyclopedias represent a microcosm of the world's knowledge. Their power comes from filtering and editing knowledge in order to highlight what is important and providing consistent coverage and integration across broad areas of knowledge. This power comes at a price, for no one encyclopedia is appropriate for readers of varying backgrounds, abilities, and purposes. Also, the bulky printed form hinders access to the rich organization and interconnections among articles.

This research was sponsored in part by Atari, Incorporated, in part by the National Science Foundation under grant MCS-8202520, and in part by a grant from Atari to the Computer Science Department at the University of Washington. Hewlett-Packard provided additional support for writing and publishing this article.
Authors' present addresses: S. A. Weyer, Hewlett-Packard Laboratories, 1501 Page Mill Road, 3U, Palo Alto, CA 94304. Weyer.hplabs@csnet-relay. A. H. Borning, Department of Computer Science, FR-35, University of Washington, Seattle, WA 98195; Borning@washington.

Personalization suffers too because of our reluctance to annotate the text and margins of articles with questions, insights, and editorial suggestions—annotations that would be impractical to suppress, share, or move.

At present, commercially available electronic encyclopedias, such as *Encyclopaedia Britannica* [20] on Mead Data Central's NEXIS and *Academic American* [1] on Compuserve, Dow Jones, and several dial-up and videotex services [8], offer rapid random access, some pattern search capabilities, and the display of article structure. For reading, paper is to be preferred over user interfaces with small text-only displays that inflict a sense of "tunnel vision" and limited context. Free-text querying (as in the NEXIS system, where any word or word pattern occurring in an article can be located) provides a useful adjunct to traditional indexes, mainly for unique terms and proper names. It is less successful, however, for words with many meanings or for specifying concepts that may be expressed in many ways or that include subconcepts.

Although the focus of this paper is on encyclopedias, many of the techniques described here have additional applications that involve the need to access rich and interconnected knowledge bases and that may include pictures and dynamic objects, as well as text. Examples of other possible applications include decision support systems, on-line manuals and other office information systems, and electronic textbooks.

2. A VISION OF INTEGRATED, POWERFUL KNOWLEDGE RESOURCES

The implementation described in this paper is an initial step toward a much more ambitious vision of an electronic encyclopedia. To provide a context for what follows, it is useful to describe briefy how the encyclopedia that we envision will eventually be written. (A longer description of this vision may be found in [16].) This perspective should apply to a larger *electronic community of knowledge* that includes many forms of on-line information, for example, newspapers, journals, messages, notebooks, and discussions. We consider encyclopedias to be central to such a knowledge network.

The information in this future encyclopedia should be as comprehensive and detailed as in the best current print encyclopedias. However, full advantage should be taken of the possibilities of the new medium, which should include not just text and static pictures, but also video sequences, animation, simulations, music, and voice. Rather than have information stored as fixed text, we envision its being encoded primarily in a concept network (making use of work in artificial intelligence on knowledge representation), augmented with literal text (e.g., of a poem), pictures, simulations, and so forth. In general, what the user sees would be custom-generated and based on the encyclopedia system's model of the user's interests, vocabulary, knowledge of specific subjects, and previous interactions with the system. In general, text would be generated automatically from the concept network using techniques such as those employed in [18]; simulations and video sequences would be custom-selected for the given user.

The idea of browsing has strongly influenced our present and future approach toward interacting with knowledge. Perusing the stacks of a small library and leafing through a magazine are common examples of browsing: navigating through a neighborhood of information and referencing items by pointing or recognizing, in contrast to accessing items directly using a known name (as through a subject index or card catalog). Browsing operates over a map or organization of the information, so it is still structured. Maps can become large, however, and difficult to use for navigation. We choose to think of the items in a browser as providing choices for the user, but as being essentially passive.

As more knowledge about the user becomes available to a future encyclopedia system, the encyclopedia should be able to recommend choices regarding paths through the information and the amount of detail to reveal in the map and underlying subject, and more personal choices with respect to style of presentation. In moving beyond browsing, we have chosen four metaphors to describe the aspects of a user's session with the encyclopedia: *models*, *tours*, *filters*, and *guides*.

2.1 Beyond Browsing: Four Metaphors

A *model* is a representation of some knowledge. As previously discussed, it may be encoded in a variety of ways.

A *tour* is a particular path through some information (the model). It can be large, for example, a presentation of the history of Britain, or small, such as a tour describing the signing of the Magna Carta. Tours would often include choice points, at which the user could decide which direction to take next. Although the basic framework for a tour may be "canned," in general, tours would be constructed dynamically in response to the interests and desires of the user.

A *filter* is the logical analog of an optical filter. Placed between a model and an observer (the user), it can mask out detail, add emphasis, combine information from several sources, and help determine presentation style, as well as perform more mundane tasks like determining whether metric or English units should be displayed and which type font should be used.

A *guide* is the user's personal (electronic) agent in the encyclopedia system. The guide sets up tours, explains, helps select filters, points out interesting topics, and provides help when requested. The guide also builds up a description of the user's preferences to better tailor the tours that the encyclopedia provides. Guides might have different personalities and styles; the user could pick a guide according to his or her tastes, for example, the Renaissance balance of Leonardo da Vinci or the novelty of the latest rock star.

In terms of the models–tours–filter–guides metaphor, using this future encyclopedia bears little resemblance to reading through a set of tomes, but is better viewed instead as a conversation with a guide or tutor who accompanies us during our learning adventure in an electronic amusement park or interactive science museum.

2.2 Areas for Research

Before such a system can be implemented, ideas from a number of diverse research areas need to be integrated and further work must be done. Relevant research areas include

—*Knowledge representation techniques.* The state of the art of knowledge representation (see, e.g., [3] and [12]) is probably not adequate to encompass all the sorts of information found in an encyclopedia.

—*Authoring systems.* The "authors" of articles in the electronic encyclopedia will need sophisticated computer support for entering and viewing information. This support is likely to include editors for knowledge networks, systems for constructing simulations (see Section 4.3), composition tools for putting together useful video sequences and animations, natural language understanding systems to allow information in an existing print encyclopedia to be digested automatically and entered into the new system, an infrastructure that allows convenient communication with editors and collaborators at remote locations, and more.

—*Natural language generation.* As described above, in general, the encyclopedia should custom-generate the text to be displayed to the reader from its internal representation. Again, work has been done in this area (see, e.g., [18]), but more is needed.

—*Agents.* To make systems more user friendly, the user interface may evolve into an "electronic personality" (a computer persona with human qualities and idiosyncrasies). In addition to solving the technical problems of how a user could construct and modify a personalized agent, attention must be paid as well to the psychological benefits and dangers of anthropomorphizing the system.

—*Intelligent tutors.* Clearly, this future encyclopedia has great potential as a tool for education, as well as for reference, and should include tutoring capabilities. Information repositories provide storage and retrieval, but little in the way of explanation, experimentation, and discussion; they do not distinguish among data, information, knowledge, and wisdom. In contrast, a talented tutor can make a subject come alive through dialogue, which is potentially different for different individuals. A number of interesting tutoring systems have been constructed for specific domains (see, e.g., [27], a collection of papers in this area), but no system anywhere near the scope required for an encyclopedia has been attempted.

—*User modeling and customization.* Research is needed on modeling what the reader knows and has seen, the reader's preferences, likely pitfalls for a particular reader, and so on. As before, some systems with user modeling capabilities have been constructed [27], but not of the scope needed here.

3. PROTOTYPE ENCYCLOPEDIA BROWSER

The future encyclopedia is clearly an ambitious, long-range project. To start exploring some of the questions raised by such a project, we have constructed a much simpler prototype version. In this prototype, instead of scrapping text as the primary means of knowledge representation, we have taken the text from an existing encyclopedia [1] and, when possible, extracted machine-manipulable knowledge from the textual representation. For selected articles we have gone on to experiment with other ways of presenting information: Some of these articles include interactive simulations; others have links to pictures stored on a videodisc.

This mixture of static and dynamic knowledge in its primitive form constitutes our model. There is no guide for suggesting tours—that choice is for the individual reader. We have chosen to concentrate our efforts on providing different kinds of filters.

In this section we present a tour through the prototype encyclopedia, looking first at an overall organization of the screen presentation, then at active text and filters, and finally at simulations.

3.1 Browsers and Menus

The prototype makes use of a window-oriented browser, with each browser window containing a set of nonoverlapping panes. The browser uses menus extensively for making selections by positioning the cursor with the mouse and clicking one of the buttons on the mouse. For example, Figure 1 shows the result of selecting the Africa article in the index. Although most of the figures show just one full-screen window, in fact, many encyclopedia windows can be created and selected through the window manager, as in Figure 2. These windows can be reshaped, maintaining the relative sizes and positions of the panes, so that parts of several encyclopedia browsers are visible simultaneously.

In general, an encyclopedia browser (as in Figure 1) consists of an article index (an alphabetical list of all articles, at upper left), an article list (list of selected articles chosen by the user, at upper right), a table of contents for an individual article (containing sections and subsections, at left below the index), and a text area (occupying the lower right portion of the browser). As seen in some later examples, these features may have different display manifestations and locations.

The index contains names of articles and cross-references; selecting a name with the mouse causes the corresponding article to be displayed. However, index entries are more than just an alphabetical list of article titles. In some cases they contain a hierarchical structure for an article, indicated by the symbol "⟩". When the "⟩" symbol is selected following an index entry, the next level of index is inserted indented and "⟨" appears for later collapsing of this level (see Figure 3). Alphabetically organized entries at the next level might correspond to sections or subsections (e.g., "Agriculture") in the parent article, or refer to entire other articles.

The table of contents for an article contains at least the title and author. For longer articles it may also refer to major sections and subsections, "see also" references, bibliography, and tables. Selecting an entry in the table of contents displays the corresponding text on the right, for example, "Introduction" in Figure 1.

3.2 Active Text

A novel feature of our computer-based implementation is the use of active text for such purposes as browsing to cross-references, converting measurements to different units, selecting videodisc images, and expanding abbreviations. Active text objects are indicated by using a special typeface: bold for cross-references, italics for convertible measurements and picture references, and uppercase for abbreviations. Active objects are also indicated by an outlined region when the cursor is nearby. The menus described in the previous subsection are implemented in terms of this active text mechanism as well.

Cross-references are names of other articles. These are usually displayed in boldface, for example, "**Africa, History of**," and "**African Art**" in the text in Figure 1. Selecting a cross-reference with the mouse will cause it to be added to

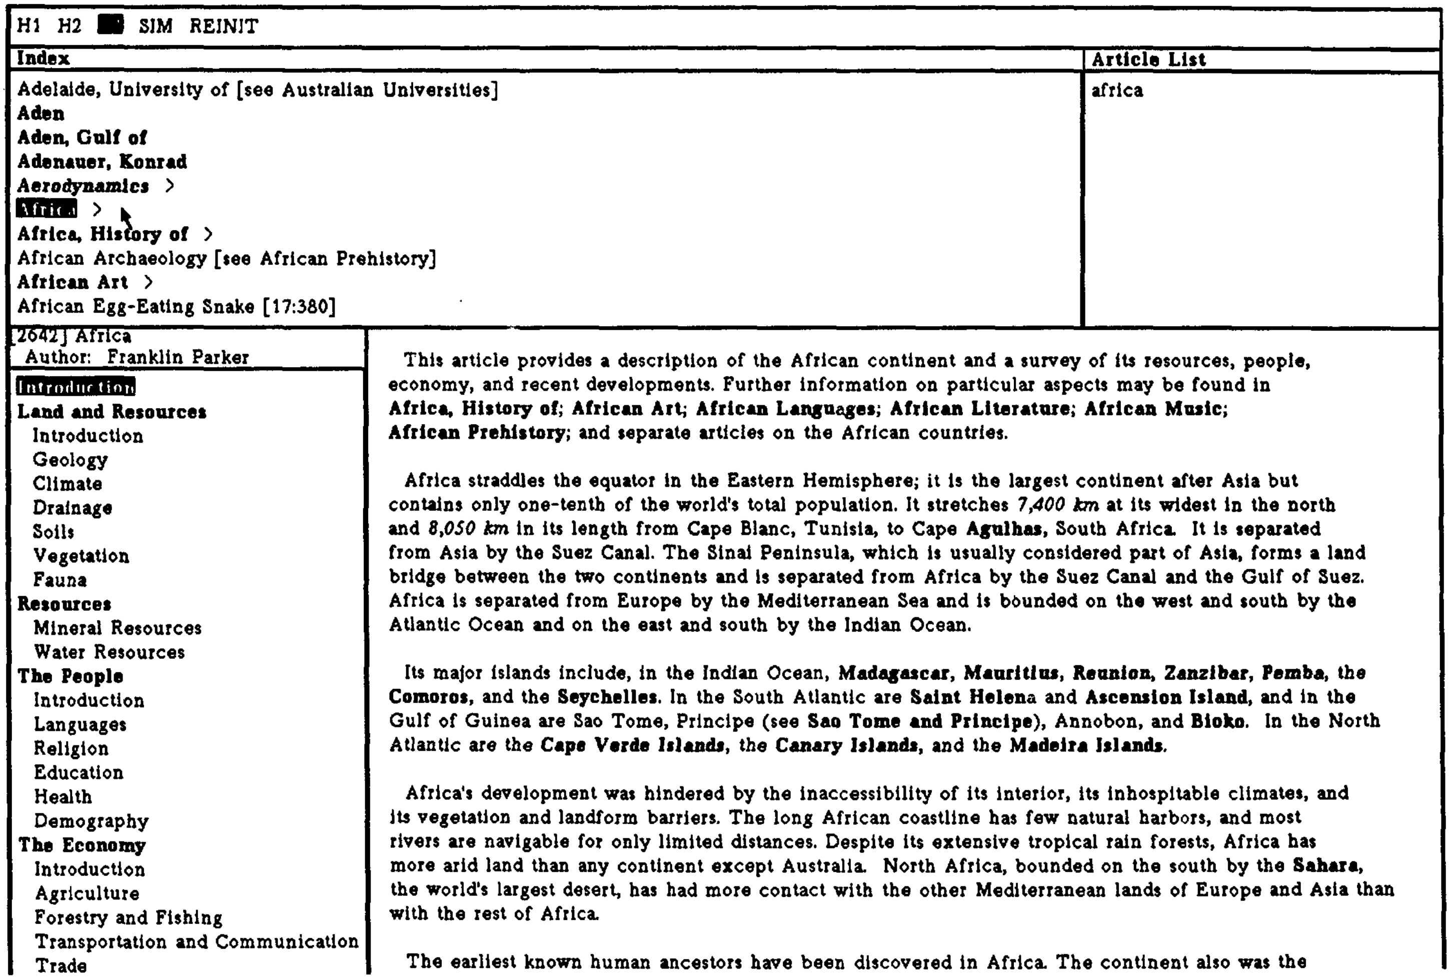

Fig. 1. Encyclopedia browser with article on Africa selected.

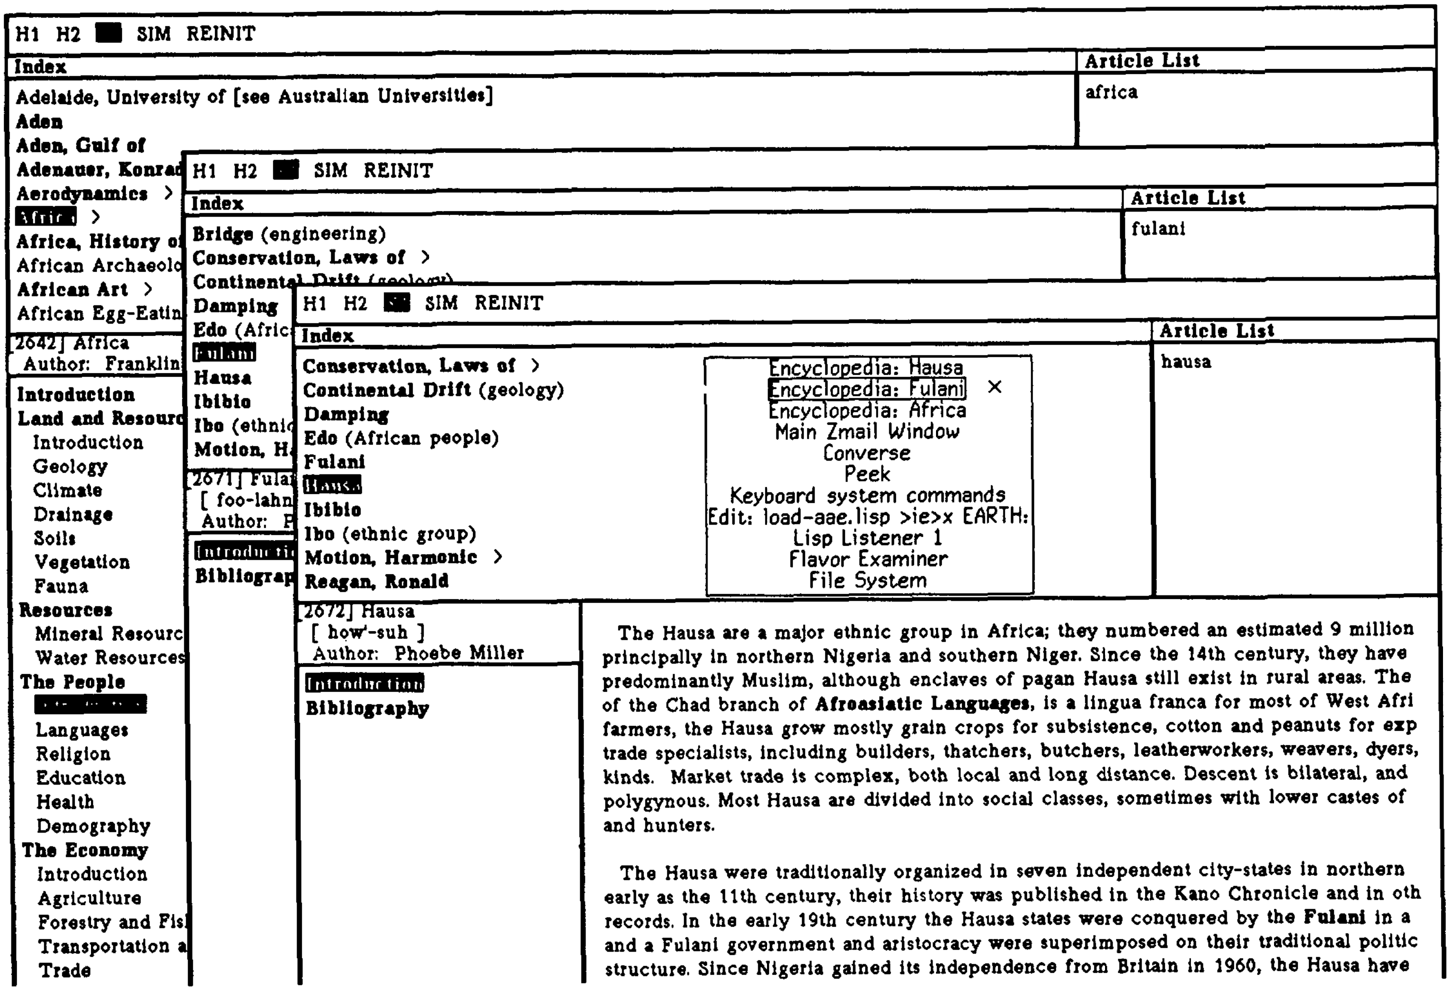

Fig. 2. Multiple encyclopedia browsers.

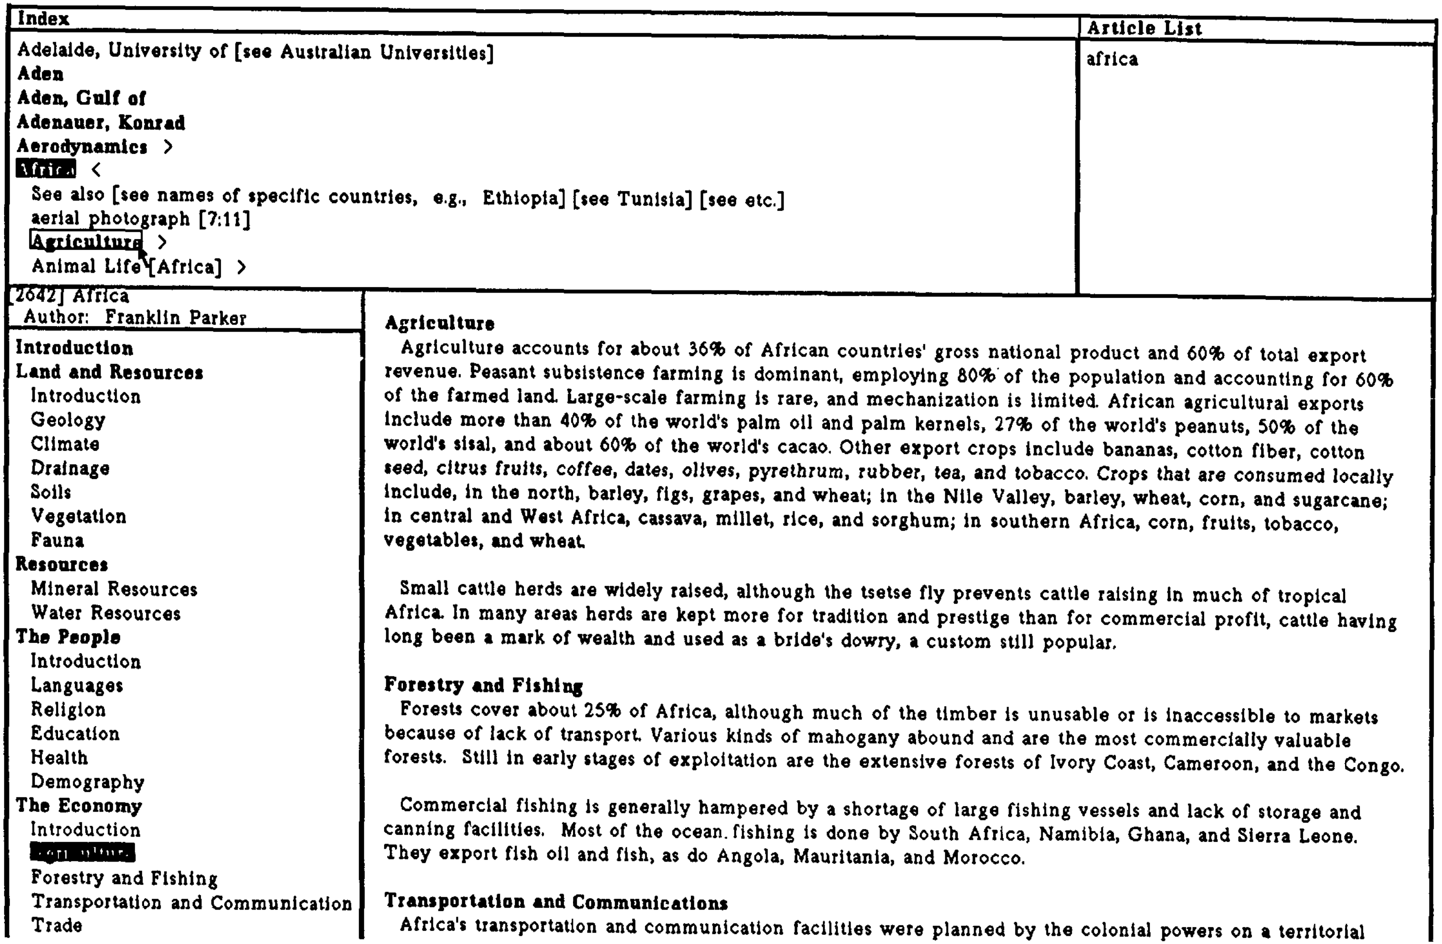

Fig. 3. Moving down a level in the index hierarchy.

the Article List for later viewing. The article list is a simple form of search history for keeping track of articles visited and cross-references yet to be followed. It allows the names of a number of articles of interest to be accumulated and supports convenient movement among several topics. An example of adding a cross-reference to the article list is shown in a later figure (Figure 10 in Section 3.4).

The distance "*5,000 mi*" in the lower half of the browser shown in Figure 4 is an active text object. When it is selected with the mouse, the metric equivalent "*8,050 km*" temporarily overlays the window. Other quantities appearing in the text, such as weights, areas, or volumes, may similarly be converted from English to metric units, or vice versa.

Tabular information (see Figure 5) appears for selected articles, for example, for articles that describe groupings of instances (bridges, military aircraft) or properties (such as for countries or presidents). Many of the measurements here are active objects as well, as indicated by the italic typeface. Currently, each entry in the table is treated only as text, and a clear direction for improvement is further decomposition of these entries into database attributes and values that can be compared and indexed.

Figure 6 shows another active text object, a picture reference. When one of the animal names in italics in the "Fauna" subsection is selected with the mouse, a picture of that animal is retrieved from a videodisc and displayed on an adjacent color monitor. Note also the alternate way of accessing images through a list of captions beginning under "Illustrations" at the end of the table of contents in Figure 5.

Abbreviations are another example of active text objects in Figure 7. "EEC" is an article-specific abbreviation recognized by the article parser. Selecting an instance of EEC provides its expansion into European Economic Community in much the same way as equivalent measurements were displayed. This same technique is used for global abbreviations and could be extended to describe any word found in a companion on-line dictionary.

3.3 Filters

In the vision of the future described in Section 2, we discussed the use of filters for such tasks as masking out detail, adding emphasis, combining information from several sources, selecting presentation format, or determining whether metric or English units should be displayed and which typeface should be used. Our prototype implementation includes some elementary filtering capabilities.

One kind of filtering is used to change the window layout by selecting one of several fixed configurations (shown in the top line of the browser); these control the size and location of the panes and the format of the information contained within the panes. One configuration is as shown in Figure 1; another configuration (not shown) uses a wider text area at the bottom with smaller typefaces, as well as two side-by-side panes for sections and subsections, compared with the single indented table of contents on the left in Figure 1. An example of a third configuration is that used in simulations (e.g., Figure 9).

Two other forms of filtering are also demonstrated in Figure 1. Depending on which button on the mouse is used to select an article, section, or subsection,

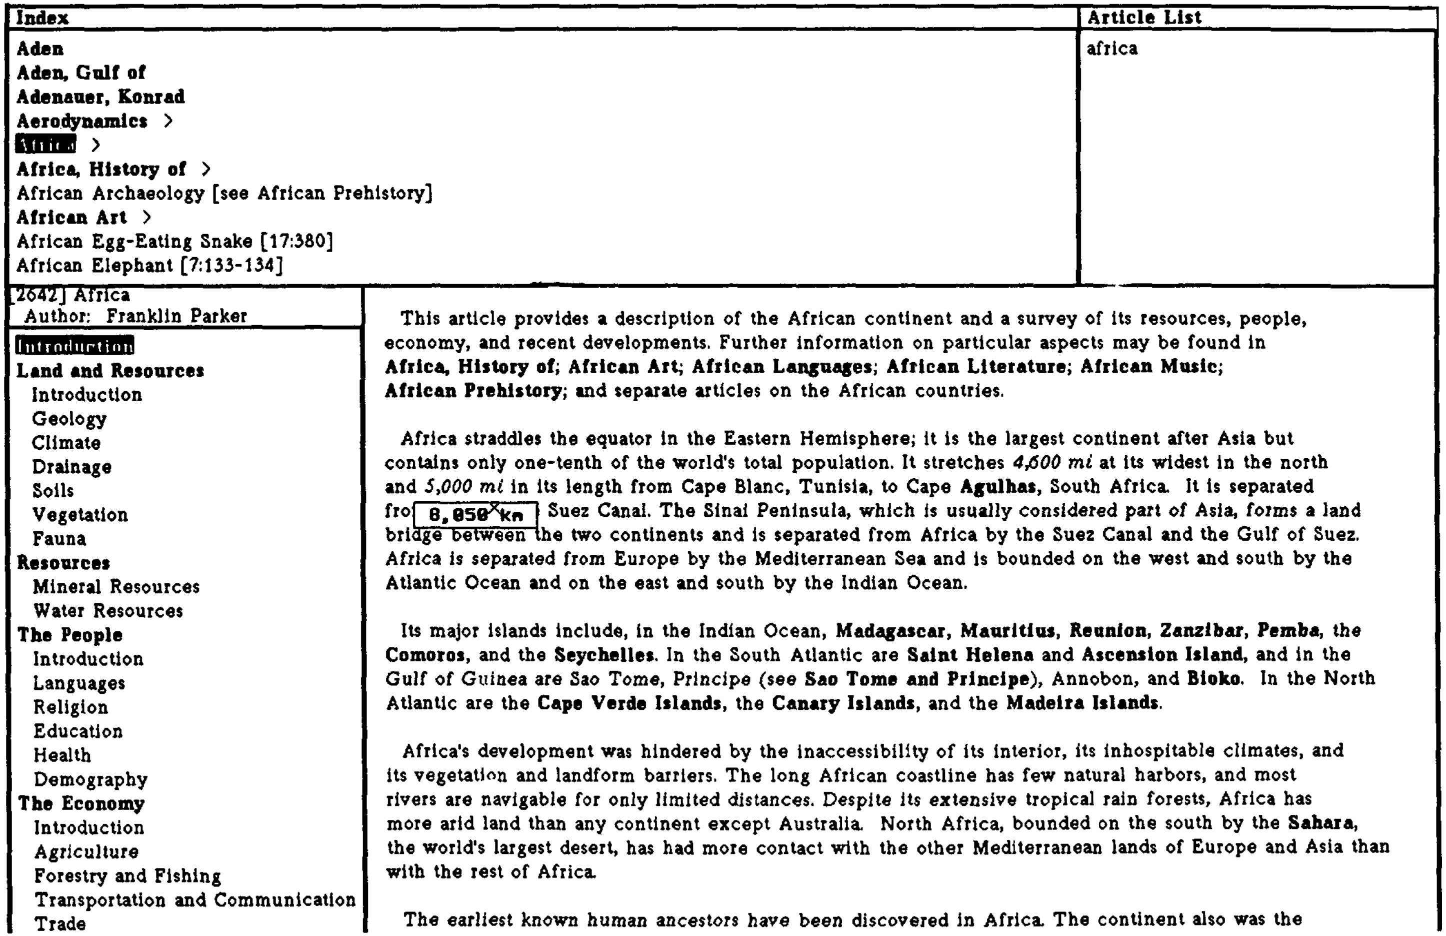

Fig. 4. Converting a measurement from English to metric units.

Index

Aerodynamics >
Africa >
Africa, History of >
African Archaeology [see African Prehistory]
African Art >
African Egg-Eating Snake [17:380]
African Elephant [7:133-134]
African Hunting Dog
African Killer Bee [3:158] [3:161-162]
African Languages [12:355] >

Article List

africa

[2642] Africa
Author: Franklin Parker

Forestry and Fishing
Transportation and Communication
Trade
Recent developments
Bibliography
GENERAL
GEOGRAPHY
ECONOMICS
PEOPLES AND CULTURES
PLANT AND ANIMAL LIFE
Fact Box 1
Area
Population
Coastline
Elevation
Principal Rivers
Principal Lakes
Principal Mountain Ranges
Principal Deserts
Political Divisions
Largest City
Busiest Port
Most Populous Country
Illustrations
anteater

Area
30,300,000 sq km;

25% of the world's land area

Population
455,000,000 (1978 est.);

10% of the total world population;

Density--*15 persons per sq km;*

annual rate of increase, 2.7%

Coastline
30,500 km

Elevation
Highest--Kilimanjaro, *5,895 m;*

lowest--Lake Assal, *155 m* below sea level

Principal Rivers
Nile, Congo (Zaire), Niger, Zambezi, Orange, Volta

Principal Lakes

Fig. 5. A table summarizing information about Africa.

Index

Adelaide, University of [see Australian Universities]
Aden
Aden, Gulf of
Adenauer, Konrad
Aerodynamics >
Africa >
Africa, History of >
African Archaeology [see African Prehistory]
African Art >
African Egg-Eating Snake [17:380]

Article List

africa

[2642] Africa
Author: Franklin Parker

Introduction
Land and Resources
- Introduction
- Geology
- Climate
- Drainage
- Soils
- Vegetation
- Fauna

Resources
- Mineral Resources
- Water Resources

The People
- Introduction
- Languages
- Religion
- Education
- Health
- Demography

The Economy
- Introduction
- Agriculture
- Forestry and Fishing
- Transportation and Communication
- Trade

Fauna

Africa's grasslands contain most of the world's last remaining herds of wild animals, such as antelope, *giraffes*, *zebras*, and the animals that prey on them: hyenas, jackals, leopards, and lions. *Elephants*, which once freely roamed the continent but were killed by the thousands for their ivory tusks, are now common only in eastern and southeastern Africa. Rain forests contain chimpanzees, *gorillas*, monkeys, *buffalo*, and wild pigs. Baboons are common, especially in the grasslands. Tropical rivers and swamps contain *crocodiles*, *hippopotamuses*, snakes, lizards, toads, and frogs. The rock python, the largest snake in Africa, reaches a length of *8 m*. Cameroon's giant frog, the Rana goliath, is the largest in the world, weighing up to *4.5 kg*.

Africa has about 2,300 species of birds. Water birds such as flamingos, pelicans, and storks live in eastern Africa. Ostriches live in southern and eastern Africa and in the western Sahara. Other tropical birds include varieties of pheasant, dove, kingfisher, hornbill, owl, parrot, duck, goose, heron, warbler, sunbird, weaverbird, and hawk. Open grassland birds include the eagle, bustard, falcon, and quail.

Debilitating insects include various mosquitoes, which carry malaria and yellow fever, and the tsetse fly, which carries sleeping sickness. Poisonous snakes of the rain forest include the cobra, vipers, and the black mamba.

South Africa's **Kruger National Park** (1898) covers *21,000 sq km*. Africa's other major game reserves include Tanzania's **Serengeti National Park**, Kenya's Marsabit Nature Reserve, and Zambia's Kafue National Park.

Fig. 6. Selecting a picture reference.

Index

Third World [19:170]
trade
Transportation and Communication >
Vegetation >
Africa, History of >
African Archaeology [see African Prehistory]
African Art >
African Egg-Eating Snake [17:380]
African Elephant [7:133-134]
African Hunting Dog

Article List

africa

[2642] Africa
Author: Franklin Parker

Introduction
Land and Resources
Introduction
Geology
Climate
Drainage
Soils
Vegetation
Fauna
Resources
Mineral Resources
Water Resources
The People
Introduction
Languages
Religion
Education
Health
Demography
The Economy
Introduction
Agriculture
Forestry and Fishing
Transportation and Communication
Trade

Trade

Africa's share of world trade, excluding gold, fell from 4.8% of world exports in 1972 to 4.2% in 1978; imports remained fairly constant at 4.5% during the same period. The oil-exporting nations, however, had a more favorable trade position. While Africa's output of crops and minerals increased in the 1950s and '60s, income did not increase proportionately because world prices declined for primary products. Single-crop export economies are still common: peanuts in Senegal and Gambia; cacao in Ghana; cotton in Sudan; and coffee in Ethiopia.

The Lome Convention, concluded in 1975 in Lome, Togo, joined 46 African, Caribbean, and Pacific (ACP) countries in trade cooperation with the European Economic Community (EEC) and replaced the earlier Yaounde Convention (between 19 African countries and the EEC). ACP countries (mainly African) sent about half of their exports to the EEC and took about half of their imports from the EEC. The United States and Canada s **European Economic Community** orts and buy a slightly higher percentage of Africa's exports. The amount of internal trade among African countries is small. Africa's trade is assisted by the United Nations Economic Commission for Africa, which began operating the African Development Bank in 1966.

Fig. 7. Another active text object—an abbreviation.

active objects will or will not be highlighted in the text. Another filter controls whether to display measurements initially in the main text in metric or English units. In general, the default format for measurements is metric; in Figure 4, the format has been changed to English. Displaying active objects themselves can be thought of as using a conversion or expansion filter.

Finally, filters are used to control how much text is to be displayed when a subsection is selected. One option is to display the subsection, followed by as much of any subsections immediately following it as will fit in the window. Alternate settings for the filter are the display of just a single subsection (and none of its neighbors), or the display of the entire article as a linear collection of headings and paragraphs. When multiple subsections are displayed, the subsection headings are also active objects, which can be selected and scrolled to the top of the text area.

3.4 Simulations

In a number of articles, the textual description is augmented by simulations that allow the user to experiment with the topic being discussed. The abacus article (Figure 8) includes an interactive abacus (a demonstration program from Symbolics, Inc.), which is activated by selecting it in the "Simulation" section in the table of contents. Abacus beads can be moved with the mouse; keyboard commands are used to reveal the state of the abacus and perform arithmetic operations. In the figure, we moved up the "1" bead in the leftmost column, typed "=" to see the value and "−1" to perform a subtraction, watched the calculation and carries being executed, and saw the result both graphically and textually.

In a simulation in the bridge article (Figure 9), after specifying a load on the bridge, we see the effects of forces transmitted to various truss members (positive force numbers represent tension, negative numbers compression). The bridge article might naturally lead us to the spectacular collapse of the Tacoma Narrows Bridge (footage of the collapse and a series of experiments are available on videodisc). Reading about the cause of the collapse leads, in turn, to the topic of damping and harmonic motion.

Before showing the simulations in the damping article (Figure 10), we briefly demonstrate the use of cross-references.The article list already contained names from earlier examples that we had typed in or selected from the index: "africa," "abacus," and "bridge." In Figure 10, the cross-reference "**Motion, Harmonic**" in the text is being selected with the mouse and added to the article list. (To allow adding several references while browsing the current article, the system does not jump immediately to the reference article.) We then select "*Motion, Harmonic*" in the article list and move to that article.

We select "Bouncing Spring" to start the first simulation (Figure 11). A weight is attached to the end of a spring; the other end of the spring is moved in a sinusoidal motion by the driving force, indicated by the hand. The parameters of the simulation—the frequency and amplitude of the driving force, the mass of the weight, the damping, and the spring constant—can each be adjusted by selecting values on the sliders. These sliding scales are input menus that are modified by selecting the bar region and moving it up and down with the mouse. The arrow on the far left indicates the natural resonant frequency, which is computed from the other parameters. When the driving frequency is at or near

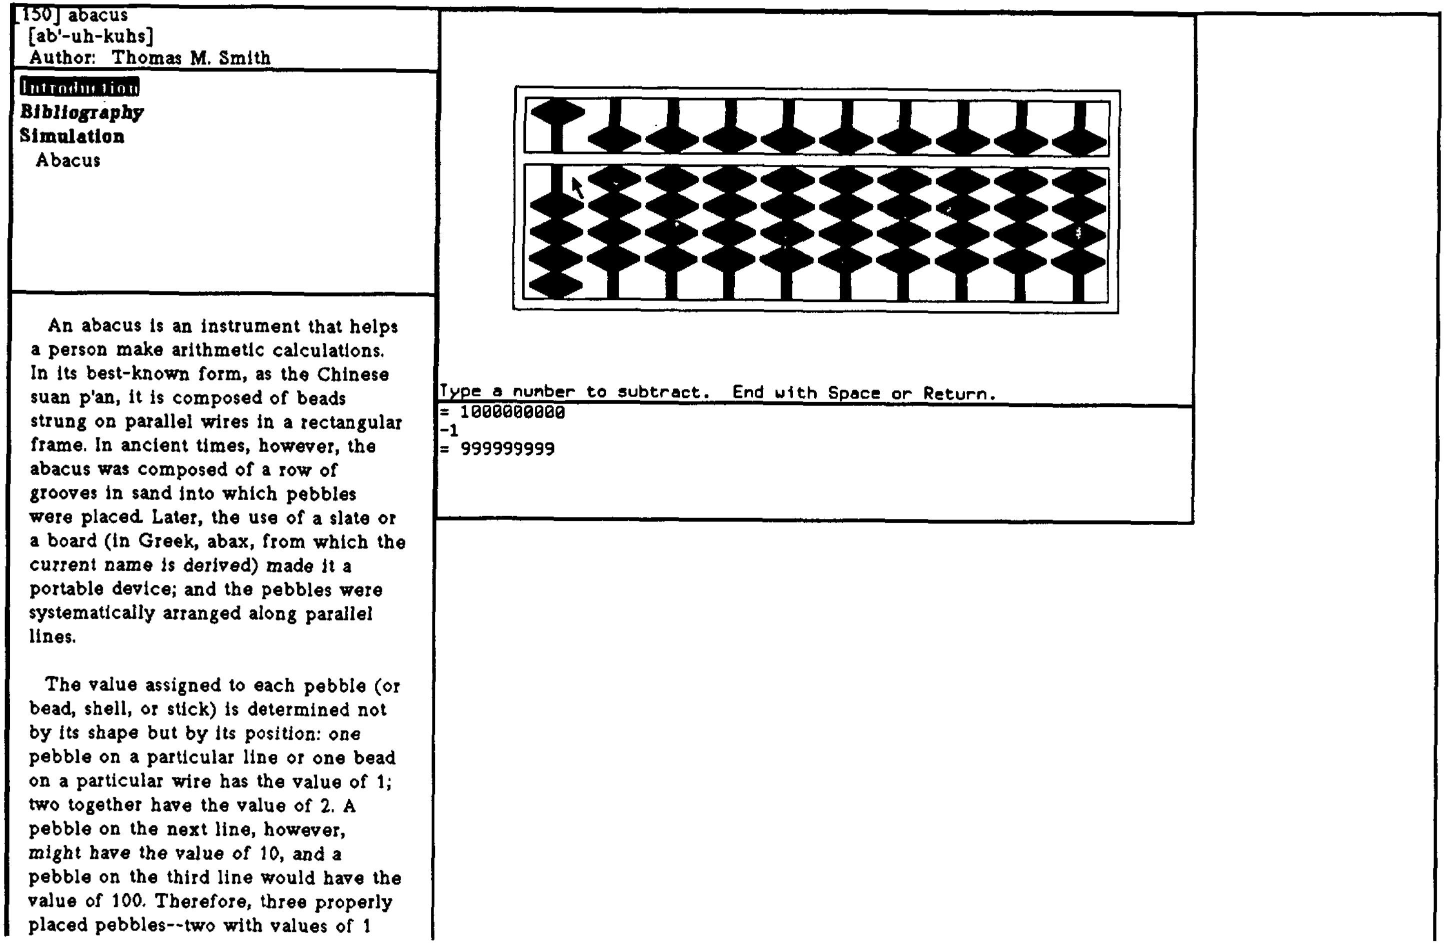

Fig. 8. An abacus simulation.

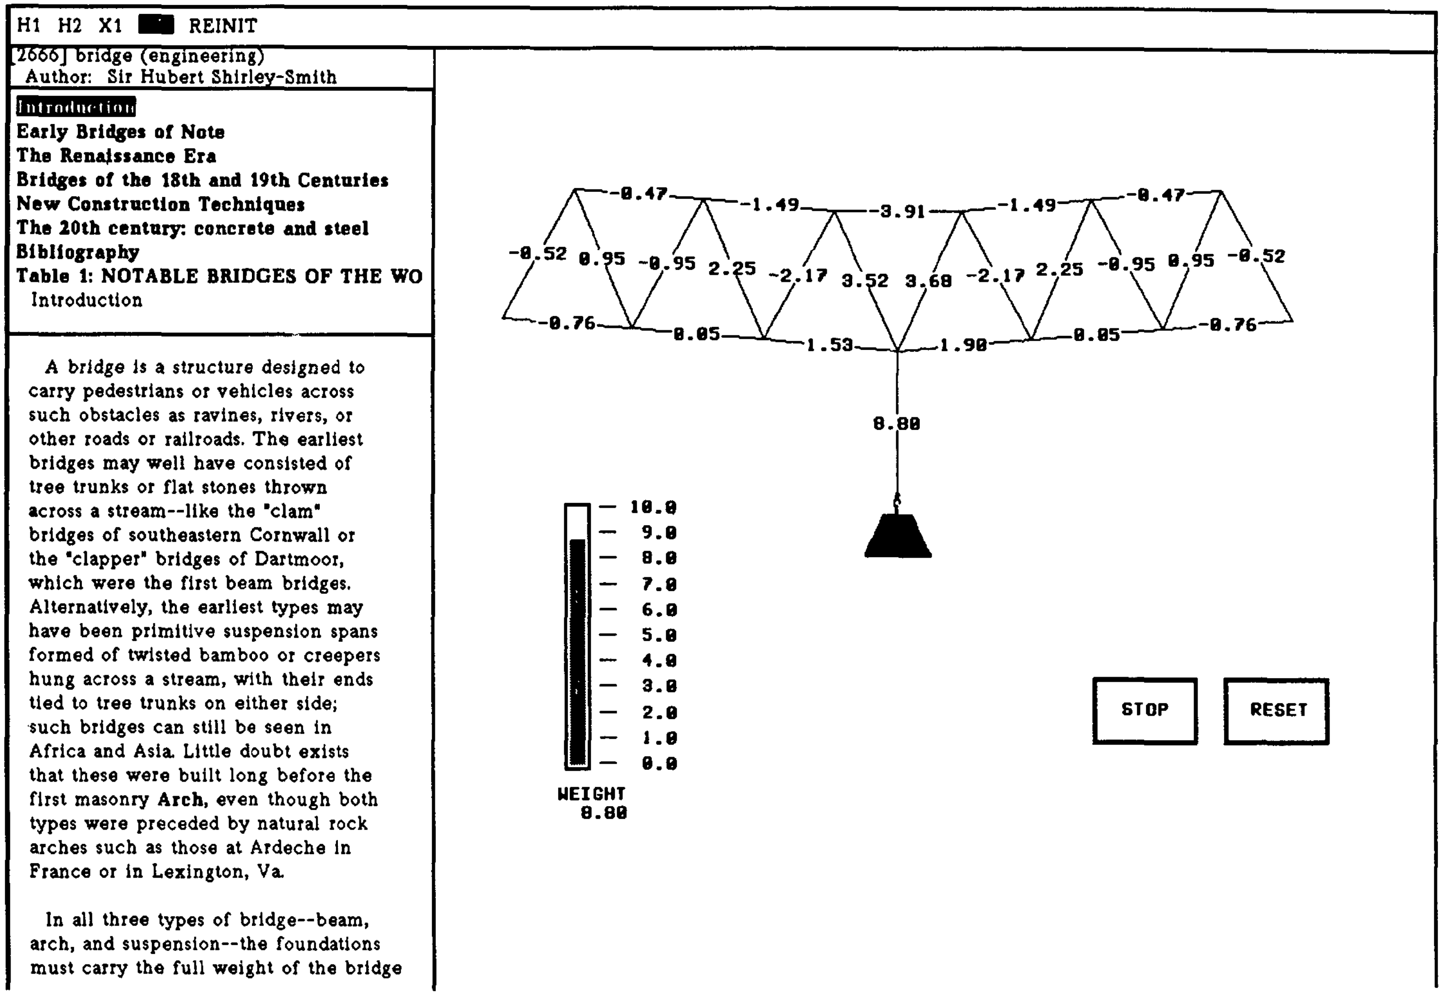

Fig. 9. Simulating a bridge under load.

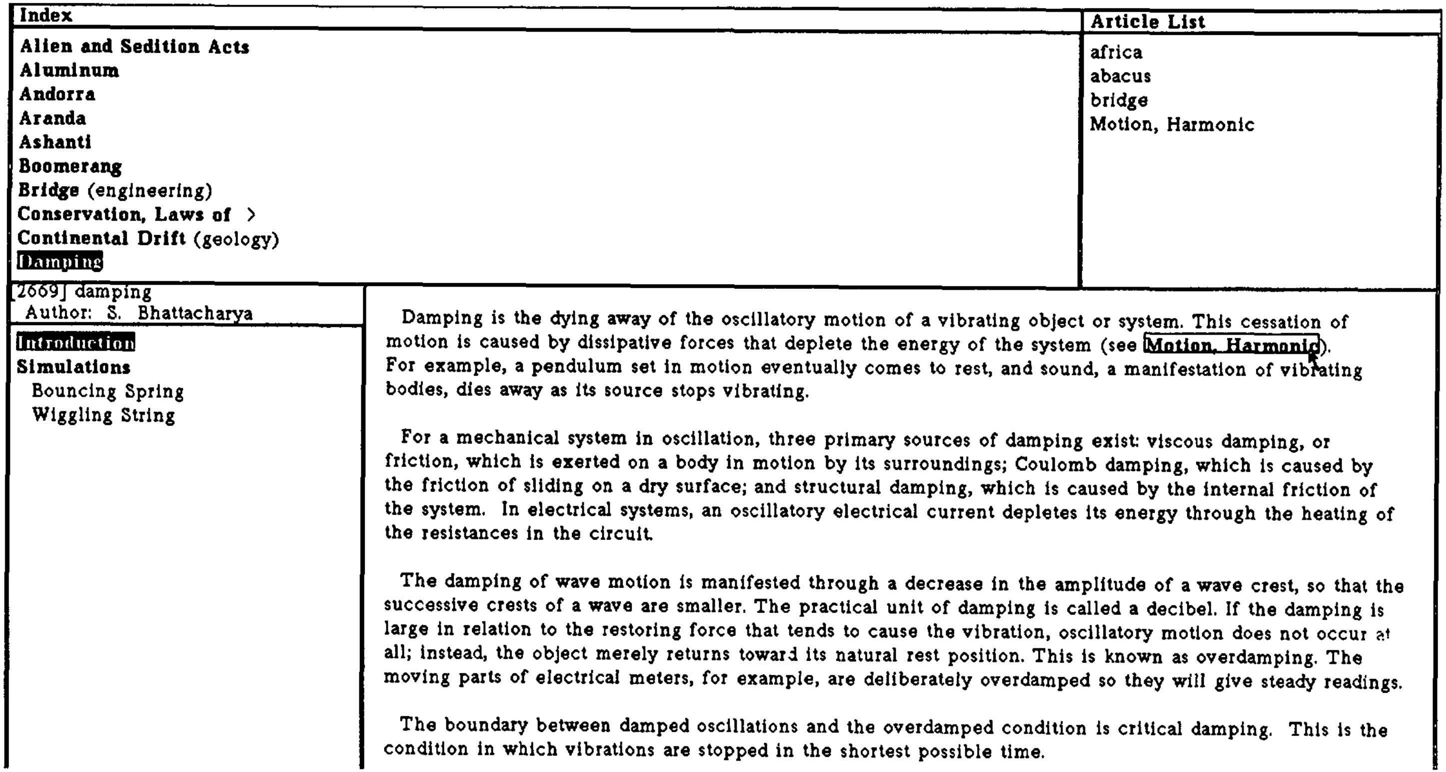

Fig. 10. The article list—an article on harmonic motion.

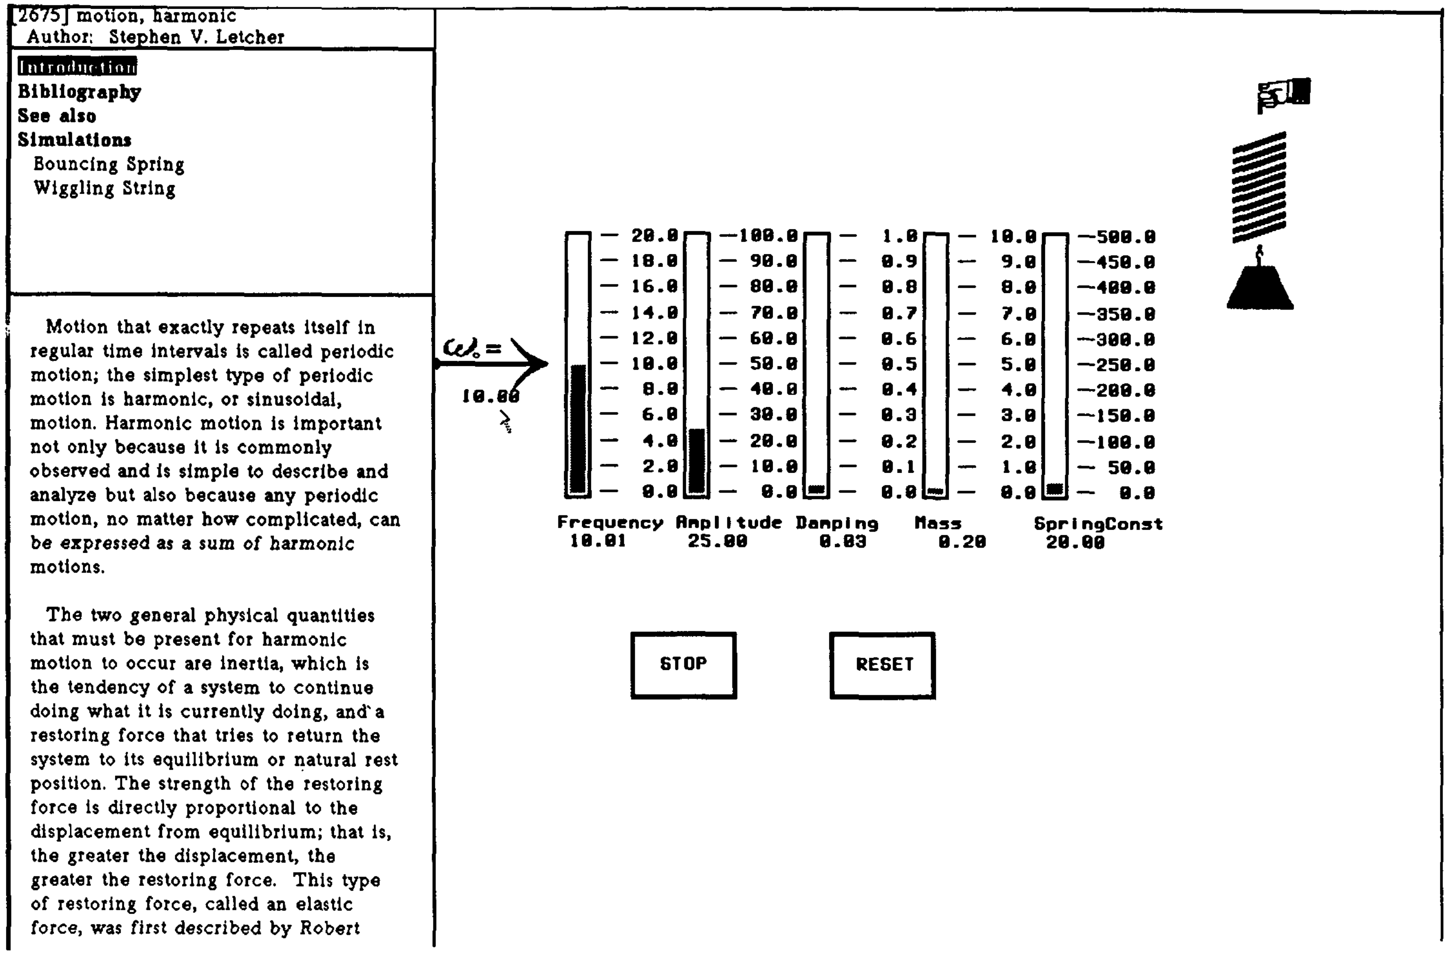

Fig. 11. A simulation of harmonic motion.

the resonant frequency, the driving force and the oscillator are coupled so that energy is continually added to the system. In the absence of sufficient dissipative forces (damping), this may lead to the result shown in Figure 12. Here, the system has detected that the spring has been compressed so tightly that it has hit the hand, and so the system has stopped the simulation and overlaid its picture with the word "SPROING!" (previously drawn using a painting program and stored as a bitmap).

A second harmonic motion simulation (Figure 13) shows a string, anchored at one end and driven by an external force at the other. (Sliders like those used in the harmonic motion simulation are used to control the vibrating string.) By varying the parameters appropriately, such phenomena as standing waves can be observed—the same phenomena that caused the collapse of the Tacoma Narrows Bridge.

4. CURRENT IMPLEMENTATION

The prototype is implemented in Zetalisp [32] and runs on Symbolics LM2 and 3600 Lisp Machines. The simulations were originally written in Interlisp-D on a Xerox Dolphin and then translated to Zetalisp; the abacus program is part of the standard Symbolics software.

In the demonstration system shown in the figures, a small collection of articles (80) and simulations (4) are contained in virtual memory. However, in the most recent version of the system, articles for the entire encyclopedia (~27000 articles)[1] have been loaded on a Lisp Machine as individuals files and can be accessed through the browser.

4.1 Articles

The implementation makes use of the object-oriented programming facilities available in Zetalisp. For example, an object (Zetalisp Flavor) known as a "p-article" (can be read as "part of article" or "particle") describes behavior that can be inherited by objects such as Encyclopedia, Article, Section, Subsection, and Paragraph. P-articles can have a heading and parts; in addition, they can be a part of another p-article. They thus model hierarchical aspects of traditional books: Sections contain textual subsections, picture captions, and simulations; subsections contain paragraphs and other kinds of objects, such as bibliography entries and table entries; and so forth. Similarly, the browser windows, both those primarily menulike and those more textual in behavior (including active objects), are defined in terms of a flavor that builds from Lisp Machine window flavors for mouse-sensitive objects, process control, text scrolling, and formatting.

From original text files supplied by Grolier's, an article parser recognizes article structure by the format of headings and builds a hierarchical description consisting of nested p-article objects. As paragraphs are scanned, article cross-references and abbreviations, both indicated in the source by all uppercase characters, are disambiguated, and active objects for these text ranges are created. Measurements matching standard formats, for example, "5,000 mi (8,050 km)," are recognized and marked.

[1] The articles from *Academic American Encyclopedia* [1] are used with permission of Grolier Electronic Publishing.

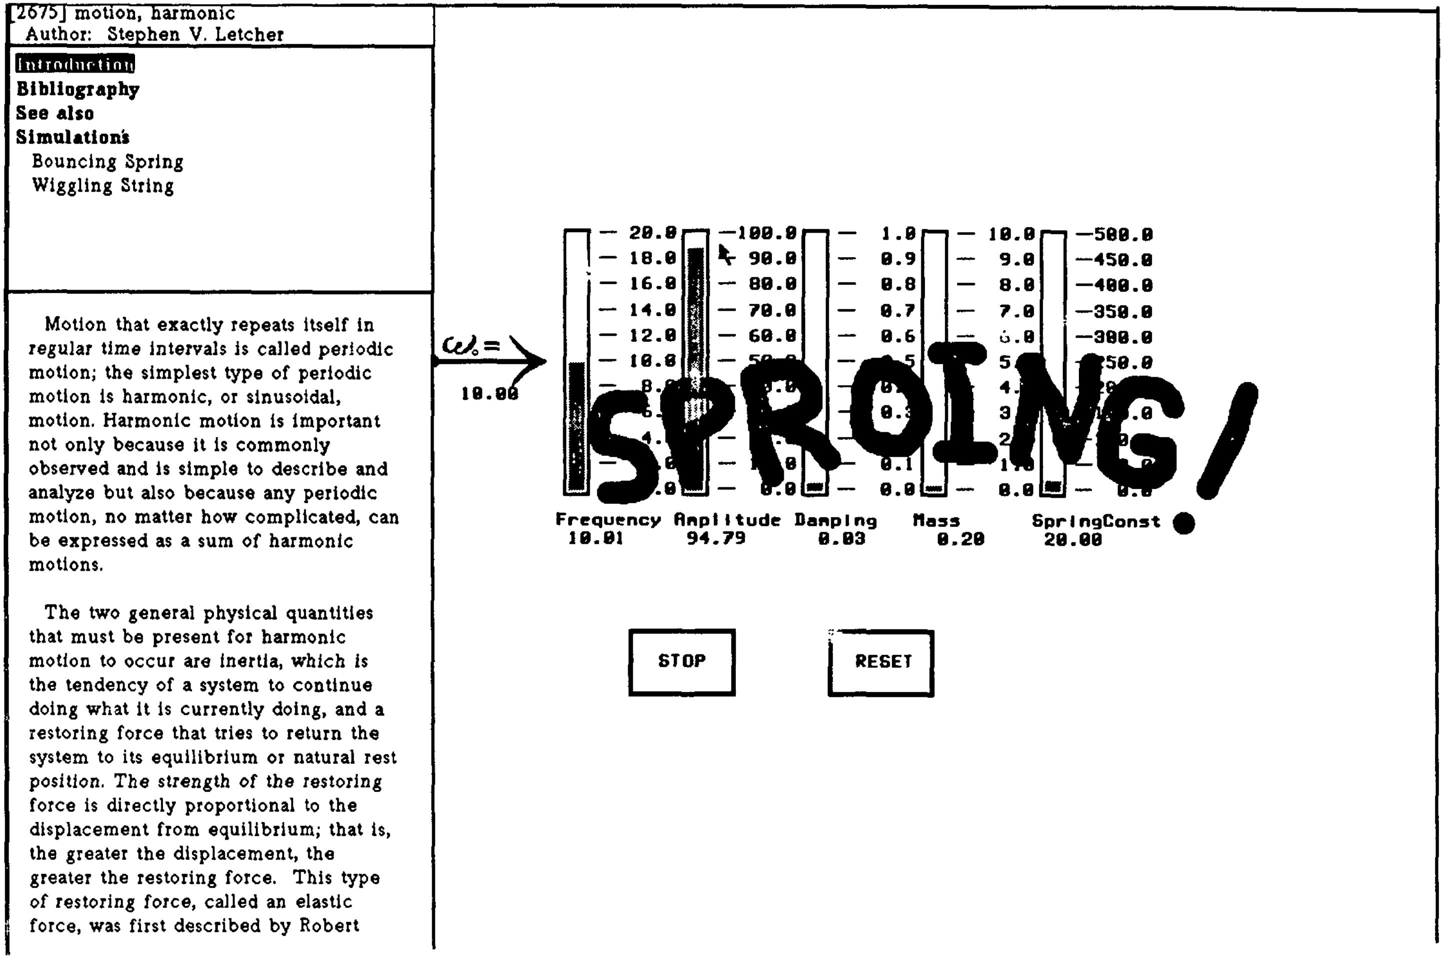

Fig. 12. The result of driving the spring too hard.

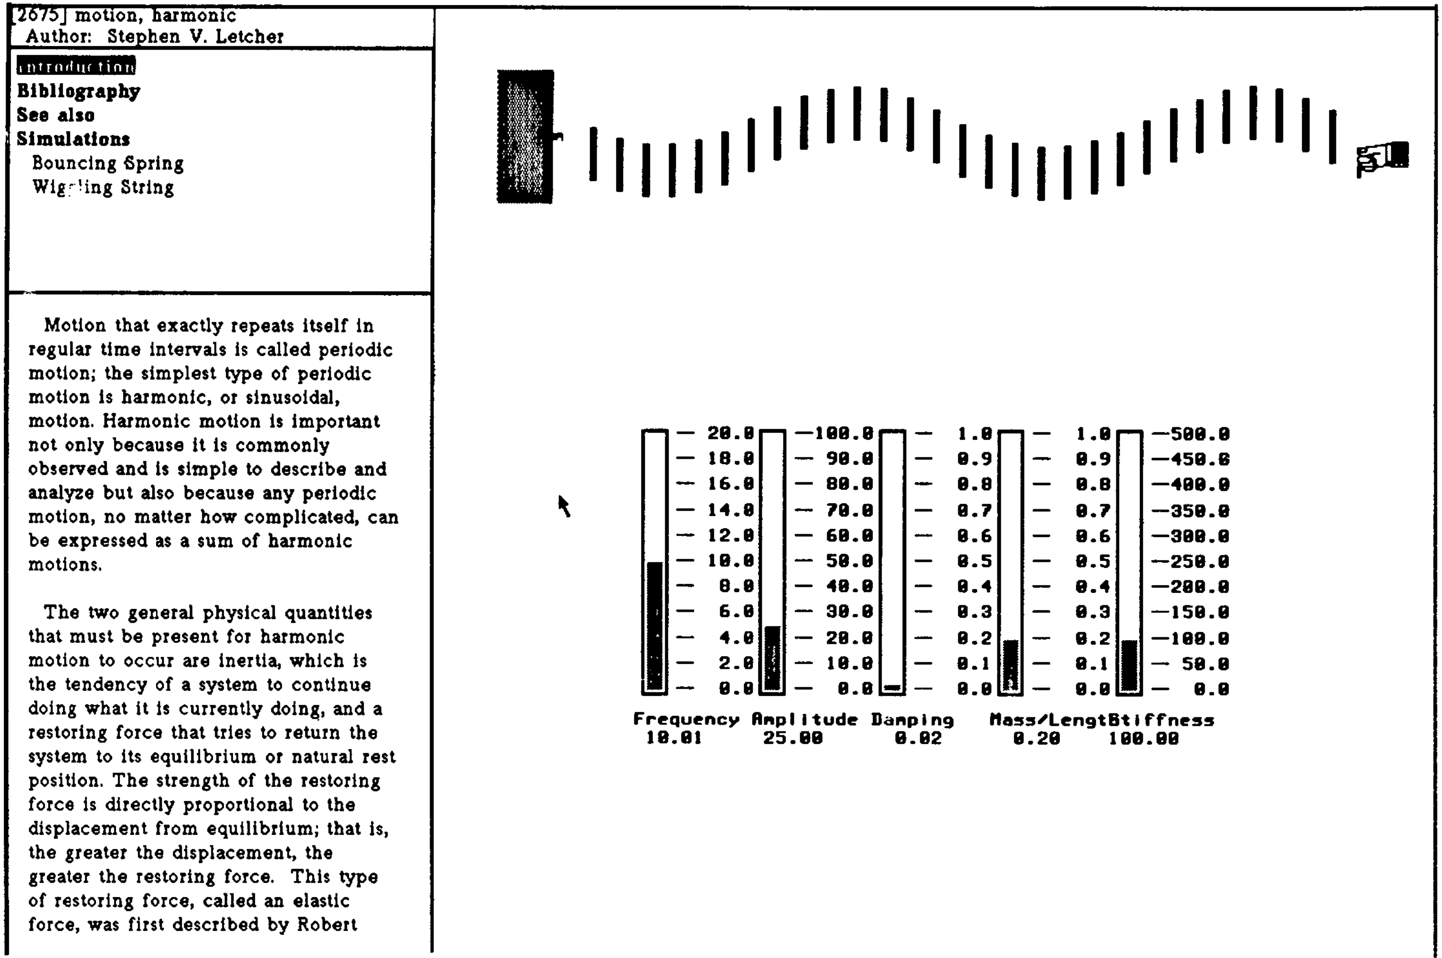

Fig. 13. A simulation of a vibrating string.

In order to add videodisc and simulation information to the articles, the original text articles are not modified. Instead, various additions are dynamically made to specific articles after they are loaded into memory. For videodisc images, a dictionary representing phrases and associated videodisc frame numbers matches against phrases occurring in running text in preselected articles. A matching phrase then becomes an active object, which is displayed in italics and the selection of which causes the associated videodisc frame to be displayed on an adjacent color monitor.

4.2 Simulations

Our simulations are more complicated and are written directly in Zetalisp. The approach taken toward merging text articles and simulations is quite primitive: A simulation is added to a specific article as a new section, for example, as in the articles on "Abacus" (Figure 8), "Bridge" (Figure 9), and "Motion, Harmonic" and "Damping" (Figures 11 and 13). Selecting a subsection of this simulation section causes the encyclopedia window to be reconfigured and the simulation to be run.

For demonstration purposes, such an approach is suitable for visualizing the desired effect of integrated dynamic articles. For an electronic encyclopedia with hundreds or even thousands of simulations, writing each as an individual program and inserting pointers into articles by hand would be prohibitively expensive. The model to be simulated could come from a variety of sources: a spreadsheet for a simple financial planning model, circuit descriptions from a VLSI layout program, or the terrain and aircraft characteristics of a flight simulator.

4.3 ThingLab

A constraint-based simulation kit based on ThingLab [4, 5] is currently being constructed with the requirements of an electronic encyclopedia specifically in mind. This system would allow the authors of the encyclopedia articles to construct simulations themselves in a convenient graphical manner. These same simulation facilities may also be used to include "laboratories" in selected articles, which allow the users of the encyclopedia to experiment with the phenomena about which they are reading.

In both the original and current versions of ThingLab, constraints are used to describe relations that must hold among parts of the simulation. For example, in constructing the resonance simulation shown in Figure 11, the constraints would be a Hooke's law constraint on the spring, an $F = mg$ constraint specifying the gravitational force acting on the weight, a constraint relating the y position of the hand to the current value of the simulation clock, and a series of constraints relating the heights of the bars in the sliders to the quantities they determine. Given a kit containing building blocks that hold these constraints, an author can rapidly assemble a simulation, leaving it to the underlying system to worry about the interactions of the parts and to keep all the constraints satisfied. A number of extensions to ThingLab's constraint representation and satisfaction techniques are being implemented to support simulations like those shown here, including constraints involving change over time, constraints that represent filters, and constraints on various kinds of collections.

4.4 Database

We planned to access the entire encyclopedia by distributing the articles on a remote machine and retrieving them over an Ethernet by attributes in addition to name, using a relational database server (Ingres from Relational Technology on a VAX 11/780). We created Zetalisp flavors to correspond to instances of databases, relations, tuples, and queries in Ingres. We specified relations to represent associations between p-article references and headings, semantic relationships among cross-references, geographical references by map latitude and longitude, and volume number and page range for each article (used to map the page references in the printed index to actual p-article references). Although we expected some response-time penalties due to this layering and network overhead, some of our anticipated advantages of this approach are time saved not constructing an ad hoc database system, increased retrieval power due to flexibility in specifying Ingres storage and index structures, and sharing among multiple users. By the time we received tapes for the entire corpus and had just begun building this large database, reductions in staff caused the termination of the project at Atari.

5. RELATED WORK

An exhaustive survey of the field would be enormous and would include work in information science, artificial intelligence, databases, education, and publishing. In this section we offer a brief tour through publications that we consider most related to the issues discussed in this paper.

In his classic paper of 40 years ago [6], Vannevar Bush envisions a desk-sized information tool (a “Memex”) that would be an electronic library and filing cabinet for helping scientists enter, view, classify, and interconnect information (via “trails” of associations) in what was considered even then to be a “growing mountain of research.” Arthur Clarke [7] speculates about fitting encyclopedias into shoeboxes and about the interface and communications technology. Nelson’s HyperText system [19] is a proposal for linking documents, versions, citations, and footnotes in such an information network; implementation is discussed in [11]. Kay [14] discusses the use of viewing and filters as a fundamental mechanism in a computer system. Our own vision of the future is expanded in [16].

In his humanist vision of the potential of this technology, Warren Preece, editor of the fifteenth edition of the *Encyclopaedia Britannica* [20], writes of “an encyclopedia for the 21st century—with a new purpose and new format to serve the global village” [21]. In a second article [22], he proposes “editorial criteria, indexing principles, and a possible format—using videodiscs instead of pages.” Van Doren [31] discusses the conception and controversial aims of a great modern encyclopedia: “teaching over informing, art over reference, human over scientific-literary, the ‘curious average man’ target, reformism over the intellectual and social status quo.” Other papers in the same volume discuss the structure of knowledge and its possible forms.

The relationship between reader and book has been augmented electronically in many directions. Cook [8] describes the state of the art in electronic encyclopedias commercially available via dial-up and videotex systems; he proposes combining these approaches with videodisc technology to construct a "multi-component electronic encyclopedia." With browsers constructed for a world history textbook and students using the browsers to answer history questions, Weyer [33, 34] explores the interplay of the search process and the structure and presentation of information. The encyclopedia browser inherits much of the window organization and search style from this history book browser but, in addition, explores the area of active objects.

Gano [10] (another member of our Encyclopedia Project) looks at the evolution of the book as a form of communication and uses videodisc and computer technology in novel ways to provide a "movie manual" for transmission repair. The Newspeek system from the Architecture Machine Group at MIT automatically illustrates text-only news articles from Mead Data Central's NEXIS with videodisc images matched to the content of the article, for example, displaying a standard picture of Henry Kissinger alongside an article about his shuttle diplomacy. (Our approach to illustrating the article "Africa" is similar.) Feiner et al. [9] concentrate on high-resolution color representations and editing of documents as directed graphs of pages, for example, in illustrating repair of complex pieces of equipment. Lee et al. [15] describe a system for constructing documents through icons and forms that integrate shared information from a variety of media. Although intended for office use, this system shares many of the qualities needed by a good encyclopedia authoring system.

Herot [13] has extended the MIT Spatial Data Management System (SDMS) to provide graphical querying and presentation of a database. ZOG [17] is a general-purpose human–computer interface system that can rapidly present "frames" (screenfuls of information), with menus being used to traverse the frames. Price [23] describes a system called "Thumb" for indexing and accessing on-line technical documentation. Stonebraker et al. [29] suggest ways in which the relational database model and the Ingres system could be used to represent the structure in documents. Trigg [30] addresses wider issues of the on-line scientific community; his "Textnet" system provides organizing principles for documents and links. Rouse [25] evaluates the use of hard-copy versus computer-based aircraft operating manuals in the context of following complex procedures.

Moving in the direction of making books intelligent, Simmons [26] organizes part of the *AI Handbook* [2] as an inferential knowledge base that can be accessed through natural language. Reggia [24] describes efforts to develop a system to retrieve synthesized clinical knowledge from an intelligent textbook of neurology and to aid in decision making. Smith [28] reviews the progress and future directions of AI as applied to information systems. McDonald [18] introduces the problem of generating natural language explanations from various representations of knowledge.

6. DIRECTIONS FOR FUTURE RESEARCH

Encyclopedias provide a large and complex knowledge environment for experimenting with ideas about viewing and browsing, knowledge representation, authoring systems, interactive simulations, natural language generation, tutoring strategies, and user modeling. Our view of the future encyclopedia proposed the model–tour–filter–guide metaphor as a style of interacting with knowledge. Our prototype, although primitive by those standards, does illustrate novel approaches in the presentation of complex knowledge.

There are several possible ways in which the future encyclopedia can be developed. One way would be to aim immediately for an encyclopedia whose information is represented completely in concept networks, simulations, and so forth. However, we believe that an evolution from primarily text-based systems, like the prototype described in this paper, may be more tractable. Steps in the evolution would include embedding many more simulations, video images, and so forth, in existing articles. Also, more and more elaborate knowledge representation structures could be used to index and interconnect the text. A convenient initial source of such information would be the structured parts of an existing database, including tables of contents, headings and cross-references within articles, latitudes and longitudes in geographic descriptions, dates in biographical and historical articles, tables, and so forth. Later, additional machine-manipulable structure could be added by the authors. Eventually, in the evolution of the electronic encyclopedia, there would be a major shift to a concept network representation, but following the evolutionary strategy, there would already be a tested structure in which to organize this information. By the time this shift occurred, the state of the art in natural language understanding might be such that much of the concept network could be produced automatically by a language understanding program from an existing corpus.

ACKNOWLEDGMENTS

Among the people who have helped with this project are the other members and consultants of the Encyclopedia Project at Atari: Steve Gano, Doug Lenat, David McDonald, Bob Stein, Craig Taylor, and Charles van Doren; Rob Duisberg and other members of the Electronic Encyclopedia Group at the University of Washington; Peter Cook at Grolier's; and Kristina Hooper and Alan Kay at Atari.

REFERENCES

1. *Academic American Encyclopedia.* Grolier, Danbury, Conn., 1983.
2. BARR, A., AND FEIGENBAUM, E., EDS. *The Handbook of Artificial Intelligence.* William Kaufmann, Los Altos, Calif., 1981.
3. BOBROW, D., AND WINOGRAD, T. An overview of KRL, a knowledge representation language. *Cognitive Sci. 1*, 1 (Jan. 1977), 3–46.
4. BORNING, A. H. ThingLab—A constraint-oriented simulation laboratory. Ph.D. dissertation, Dept. of Computer Science, Stanford Univ., Stanford, Calif., Mar. 1979. A revised version is available as Tech. Rep. SSL-79-3, Xerox Palo Alto Research Center, Palo Alto, Calif., July 1979.
5. BORNING, A. H. The programming language aspects of ThingLab, a constraint-oriented simulation laboratory. *ACM Trans. Program. Lang. Syst. 3*, 4 (Oct. 1981), 353–387.
6. BUSH, V. As we may think. *Atlantic Monthly 176* (July 1945), 101–108.
7. CLARKE, A. C. The second century of the telephone. In *The View from Serendip.* Random House, New York, 1977, pp. 215–236.
8. COOK, P. R. Electronic encyclopedias. *Byte 9*, 7 (July 1984), 151–170.
9. FEINER, S., NAGY, S., AND VAN DAM, A. An experimental system for creating and presenting interactive graphical documents. *ACM Trans. Graph. 1*, 1 (Jan. 1982), 59–77.
10. GANO, S. Forms for Electronic Books. Master's thesis, M.I.T., Cambridge, Mass., June 1983.

11. GREGORY, R. Xanadu: Hypertext from the future. *Dr. Dobb's J. 75* (Jan. 1983), 28–35.
12. GREINER, R., AND LENAT, D. B. A representation language language. In *Proceedings of the National Conference on Artificial Intelligence* (Stanford, Calif., Aug.). American Association for Artificial Intelligence, Menlo Park, Calif., 1980, pp. 165–169.
13. HEROT, C. F. Spatial management of data. *ACM Trans. Database Syst. 5*, 4 (Dec. 1980), 493–514.
14. KAY, A. C. New directions for novice programming in the 1980s. In *Programming Technology*, P. J. L. Wallis, Ed. Pergammon Infotech, Elmsford, N.Y., 1983, pp. 209–247.
15. LEE, A., WOO, C. C., AND LOCHOVSKY, F. H. Officeaid: An integrated document management system. In *Proceedings of ACM SIGOA Conference on Office Automation* (Toronto, Canada, June 25–27). ACM, New York, 1984, pp. 170–180.
16. LENAT, D., BORNING, A., MCDONALD, D., TAYLOR, C., AND WEYER, S. Knoesphere: Building expert systems with encyclopedic knowledge. In *Proceedings of the 8th International Joint Conference on Artificial Intelligence*, (Karlsruhe, West Germany). 1983, pp. 167–169.
17. MCCRACKEN, D. L., AND AKSCYN, R. M. Experience with the ZOG human-computer interface system. *Int. J. Man-Mach. Stud. 21*, 4 (Oct. 1984), 293–310.
18. MCDONALD, D. D. Natural language generation as a computational problem: An introduction. In *Computational Theories of Discourse*, M. Brady, Ed. MIT Press, Cambridge, Mass., 1982.
19. NELSON, T. H. Literary machines. Swarthmore, Pa. Available from the author.
20. PREECE, W. E., Ed. *Encyclopaedia Britannica*, 15th ed. Encyclopaedia Britannica, Chicago, 1979.
21. PREECE, W. E. Notes toward a new encyclopedia. Part 1. *Scholarly Pub. 12*, 1 (Oct. 1980), 13–30.
22. PREECE, W. E. Notes toward a new encyclopedia. Part 2. *Scholarly Pub. 12*, 2 (Jan. 1981), 141–157.
23. PRICE, L. A. Thumb: An interactive tool for accessing and maintaining text. *IEEE Trans. Syst. Man. Cybern. SMC-12*, 2 (Mar./Apr. 1982), 155–161.
24. REGGIA, J. A., PULA, T. P., PRICE, T. R., AND PERRICONE, B. T. Towards an intelligent textbook of neurology. In *Proceedings of the 4th Annual Symposium on Computer Applications in Medical Care* (Nov.). IEEE, New York, 1980, pp. 190–199.
25. ROUSE, S. H., AND ROUSE, W. B. Computer-based manuals for procedural information. *IEEE Trans. Syst. Man Cybern. SMC-10*, 8 (1980), 90–147.
26. SIMMONS, R. F. A text knowledge base for the AI handbook. Tech. Rep. TR-83-24, Dept. of Computer Sciences, Univ. of Texas at Austin, Dec. 1983.
27. SLEEMAN, D., AND BROWN, J. S., EDS. *Intelligent Tutoring Systems*. Academic Press, London, 1982.
28. SMITH, L. C. Artificial intelligence applications in information systems. In *Annual Review of Information Science and Technology*. M. E. Williams, Ed. American Society for Information Science, Washington, D.C., 1980, pp. 67–105.
29. STONEBRAKER, M., STETTNER, H., LYNN, N., KALASH, J., AND GUTTMAN, A. Document processing in a relational database system. *ACM Trans. Office Inf. Syst. 1*, 2 (Apr. 1983), 143–158.
30. TRIGG, R. H. A network-based approach to text handling for the online scientific community. Ph.D. dissertation, Dept. of Computer Science, Univ. of Maryland, College Park, Md., Nov. 1983. Published as Maryland TR-1346.
31. VAN DOREN, C. The idea of an encyclopedia. In *The Growth of Knowledge: Readings on Organization and Retrieval of Information*, M. Kochen, Ed. Wiley, New York, 1967, pp. 66–71.
32. WEINREB, D., AND MOON, D. *The Lisp Machine Manual*. Symbolics, Inc., Cambridge, Mass., 1981.
33. WEYER, S. A. Searching for information in a dynamic book. Ph.D. dissertation, Dept. of Education, Stanford Univ., Stanford, Calif., Feb. 1982. Also available as Tech. Rep. SCG-82-1, Xerox Palo Alto Research Center, Palo Alto, Calif.
34. WEYER, S. A. The design of a dynamic book for information search. *Int. J. Man-Mach. Stud. 17*, 1 (July 1982), 87–107.