# KNOESPHERE:
## BUILDING EXPRRT SYSTEMS WITH ENCYCLOPEDIC KNOWLEDGE


Douglas B. Lcnat — Reuristic Programming Project, Stanford University, Palo Alto Ca

Alan Borning — C.S. Department, University of Wash., Seattle, Washington

David McDonald — C.S. Department, University of Mass., Amherst, Mass.

Craig Taylor — Corporate Research, Atari, Inc., Sunnyvale, Ca.

Stephen Weyer — Corporate Research, Atari, inc., Sunnyvale, La.



ABSTRACT

The Knoesphere project is an attempt to build an expert system that is encyclopedic, in the breadth of coverage of its knowledge base, and in the degree of integration of that knowledge. The primary issue is how to aid users in searching complex bodies of knowledge. Our approach is to frame the system more as a museum than a set of tomes, and to have the user take more or less guided tours of the exhibits therein. The impact of such a system on everyday life ~ entertainment and, eventually, education -- is clear. We discuss its potential for progress in AI as well: a testbed for representation, speech understanding, natural language understanding and generation, automatic story generation and animation, learning, user modelling, and planning. Having an immensely broad and moderately deep knowledge base, the system may also serve as a useful testbed for exploiting analogy and metaphor as a source of power. The work is in its early stages, hence much of what we present is the design for this system, not finished results. We do calculate the magnitude of the tasks involved in such an ambitious endeavor, and give scenarios of its use.


## 1. Introduction to the Knoesphere

This paper describes an ongoing research project whose goal is to represent a comprehensive corpus of real world knowledge (both the size and scope of The Encyclopedia Britannical [1] in a knowledge base; i.e., as a structured network of concepts, rather than as pieces of text. The purpose is to enable a user to browse through that knowledge in as effective and flexible a manner as possible. The project can be conceived as the combination of three methodologies: expert systems [2] (hitherto narrow in scope), encyclopedias and online data bases [1,3] (hitherto only lightly cross-indexed assemblages of prose), and videogames (hitherto rarely educational). We now present a scenario of a session with Knoesphere. We include it for motivation of the project, and to illustrate how the various levels of modelling and tailoring and filtering can affect what the user experiences.

It is Autumn, 1995. Sitting in our living room, we connect our home computer to the relevant encyclopedic service, and don a helmet and gloves. The helmet is a lighter, faster version of [4], with separate images for each eye, speakers for each ear, and sensors to track eye and head movements. The gloves enable our hand positions to be monitored, and provide pressure feedback.

Displayed before our eyes is a trio of three-dimensional arrays, labelled Axes, Guides, and Filters. Each array is filled with icons. We point at the Axes array, and it grows larger, as though we were approaching it. We now see that its icons symbolize various sorts of dimensions: ways of organizing knowledge: by location, by time, by degree of certainty, by principal objects involved, by philosophical orientation, by academic discipline, etc. In this context, icons have three-dimensional shapes and textures. We reach out and touch three of the icons (academic discipline, certainty, philosophic orientation), and off in a small corner we see them now labelling three orthogonal axes of a coordinate system labelled Knowledge, with tiny, illegible icons dotting that space. The Axes array has shrunk back down in size, and we now enter the Guides space. It grows, to become an array of pictures of people, arranged by their personality, profession, and teaching style. We select Jean, a blustery physicist who makes us draw our own conclusions. We notice that Jean is now standing off to one side of us, occasionally reaching over and peering at entries in the final, Fillers array. It is filled with simulated clip-on sunglasses, with labels like NoAnalogies, EmphasizeTheory, and Emphasize Names. These sunglasses filter out - or stress - various features of the tour we're about to take. We select NoMath and NoConnections, and put them on; Jean winces.

We tell Jean we're ready, and the new Knowledge array (with its axes labelled Held, Certainty, and Philos) grows to fill our entire field of vision. In front of us, all around us, are all the entries in the knowledge base, organized by these three attributes. We notice that we're positioned near Physics on one axis (Jean's specialty), at NearCertainty on another axis (Jean knows I wouldn't have selected her to guide me around speculative fiction), and at Ontology along the third axis (Jean prefers Teleology, or even Phylogeny, but knows I'm more interested in the current state of things). Moving around in this space for a couple minutes, with a running commentary from Jean, the image of Jupiter catches our eye, and we ask to spend some time finding out more about it.

The Jovian sphere opens up, and inside is what appears to be miniature amusement park. It grows larger, until it appears almost life-sized and fills our field of vision. There's an arcade with pictures of the planet, a small library (presumably of books and articles on Jupiter), a scale model of the solar system in motion (with Jupiter highlighted), and a spaceship. Jean glances wistfully-at the scale model; if we hadn't put on the NoMath filter, she'd have shown us a simulation that demonstrates how the areas swept out by Jupiter in equal time periods are equal.

We browse through the textual material for a few minutes, occasionally pointing and asking Jean a question about what we've just read. Some of the books are labelled as original sources, but each one has a duplicate that's specially tailored and filtered for us, using a model of us, the clip-ons we're wearing, and even some more subtle and dynamic cues, such as our fidgeting, what specifically we've just read or skipped, etc. Since we chose to be on the Ontological plane, there is scarcely any material about what Jupiter was thought to be by the ancients, how it formed during the creation of the solar system, etc. Jean notices we're getting bored, and suggests we go in for a first-hand look. We board the rocketship, and after a visually and aurally believable liftoff find ourselves moving quickly into Jovian space. Jean knows that we prefer to control the ship ourselves, so conveniently sits in the co-pilot's seat

One of our instruments, labelled Source, tells us that the images being generated out our viewscreen are taken from early NASA flyby mission photos. Jean answers our questions, and occasionally directs us to some interesting spots to fly over. When we insist on going down into the atmosphere, Jean scowls at a dial on the instrument panel, whose needle is dropping fast. Curious (we see the altimeter is elsewhere), we lean over and notice that dial is labelled Certainty, and the needle is moving from Fact to Fiction. Various other instruments report the chemical composition, pressure, temperature, etc. of the gases we are flying through. Suddenly, large whale-like creatures now begin to hurtle past us, and we pull a lever which returns us instandy to the Field/Certainty/Philos space. We note that we are indeed at the same Field/Philos point, but our Certainty has slipped to an extreme, namely fiction. In particular, we have emerged from a book-shaped icon labelled 2001/2010. Our original point of entry -- indeed, the entire encyclopedia of relatively certain knowledge - is now far overhead along the Certainty axis. Pushing what appears to be a button on our watch, the session ends.

## IT. Architecture of the system

It is tempting to represent the architecture of Knoesphere as a single box containing the letters KB, for knowledge base, as even the inference procedures are represented within the same formalism and hence are a real, inspectable, modifiable part of the KB. Thus, this section is really about the architecture of the KB: the sorts of knowledge it will contain, how those are organized and used.

Of course a large part of the KB represents the encyclopedic knowledge: famous battles, buildings, biographies, birds, books, etc. The original text will be completely re-represented in a structured language, from which the program can generate text, or animation, or narration, or tactile sensations. But the decision as to modality is not made in advance. Regardless of output modality, the material being presented might range from a straight-forward re-representation of entries in the KB, to a dynamically-plotted story based on such data.

The knowledge base contains frames that represent heuristic rules. These deal with building and checking user models, appropriateness of various modalities and representations (including rules for when and how to cache, or pipeline, or expectation-filter). Some heuristics are domain-independent (learning strategies), but many are domain-specific; some heuristics generate plausible suggestions, some analyze and evaluate and critique [5,6].

Something which is absent from a typical encyclopedia but must be present in the Knoesphere KB is commonsense knowledge. This includes everyday physics [7,81, models of human interactions (including conversation) [9], models of human/machine dialogues, as well as facts and heuristics about teaching, question-answering, imagery, analogy, etc.

Much of that knowledge is usable for an intelligent interface to users, though additional details about people [10] are required: types of people (models of mathematicians, mailmen, and mechanical engineers) and individual people (celebrities, specific users, archtypical group representatives). Each such model includes physical appearance, goals, state of knowledge (including notations, vocabulary), memory capabilities (information processing model of human cognition), inferencing capabilities, interests and biases, how well they learn via various sorts of teaching methods, and miscellaneous attributes (strength, intelligence, etc.) Note that for each attribute, we must record several values: the true value, the value the individual believes the true value to be, the value the individual believes s/he is projecting to others, the value they are projecting to others in truth, what they would like the true value to be, what they would like to project to others, how strongly they want those last two values to become true (or to stay true).

Given that the program has access to a very rich context in which the user expressed something (eye position, previous actions, model of what s/hc is interested in and typically asks about, where his/her finger is pointing), the tasks of understanding speech and language are constrained to the point of feasibility. The final piece of the knowledge base deals with rules and algorithms for performing these tasks.

As implied above, the control structure is one of condition-action rules, executed in a forward direction, governed by a smaller and more general set of rules (actually a subset of the first set). These rules are represented in the same language as the rest of the knowledge, and are part of the KB, hence can be viewed and tailored (perhaps automatically [6]) to the individual user.

The user selects axes, a guide, and a set of filters. All three choices serve the same function: cutting down the size of the space of information nodes to display, determining what structure to give them in the display, and deciding where in that space to start the user off initially. Often, the same constraint is available to a user in any one of the three choices he makes (e.g., s/he could enter a space where s/he chooses to be at Phylogeny along one axis; or s/he could choose a guide who always includes the origins of the things he's presenting; or s/he could clip on a filter called EmphasizeOrigins). A common preference may be for three axes that deal with subject matter, relegating all other constraints to the guide's personality and to the clip-on filters selected, 9liters can be composed. For instance, one could take the NoConnections filter, and break knowledge up into separate concepts; putting the EmphasizeNames filter over that, followed by the EmphasizeOrdering filter, would result in an alphabetized list of concept names, very much like the index of a conventional print encyclopedia.

There will be some handcrafted tours - authored, iust as any work of art -- on file and accessible to the guide; this is necessary because sometimes a sequence of exhibits is effective without anyone being able to account for why it's effective; since it can't be rederived from more elementary planning principles, it must be preserved.

The user's responses arc checked against the existing model of that individual, and discrepancies are noted (and will eventually be reflected by changes in the model). This sort of learning applies to models of groups of people as well; thus, if mathematicians become less averse to examples over a period of years, the Mathematician model would slowly change to reflect that.

## III. Feasibility, Current Status, Future AI Issues

The project described in this paper is merely the expert systems "view" of an even larger effort, the design of a large "electronic community" of services and facilities, that is a focus of Atari's new Sunnyvale Research Center. Ignoring this larger vision, ignoring the unsolved problems with hardware and graphics and simulation-design, how feasible is the goal to represent [l]'s knowledge in an expert system?

Our efforts to date suggest that there will have to be one new frame added for each of the 300,000 paragraphs of material in [J]. Since a very similar frame will usually exist, it will take a person familiar with Knocsphere's data entry process (and having, say, a B.A. in the same field as the article) about one hour to exploit the analogy and enter the new frame; this works out to three years of work for 50 such people. Analysis of [l]'s *Macropaedia* reveals only 400 distinct types of articles that we, the authors of this article, must represent by our target date of 1990. Unfortunately, there are approximately 2000 everyday concepts for us to represent (double the number in Basic English). We intend to spend much of the corning decade in research trying to build such a core.

In short, this "back of the envelope" calculation shows the *representation* part of the project to be feasible (a decade of research, five years (and 200 man-years) of development). We will *not* be working on the problems of special hardware such as helmet and gloves; high quality dynamic computer animation; home computers as powerful as current Lisp machines; inexpensive yet suitably fast data links to central knowledge repositories; but several others at Atari and elsewhere are pursuing such technologies. The calculation, and the existing project to date, span only the task of re-representing the Knowledge in [1], translating it from text into a frame-based format that can be processed using existing expert systems techniques.

As mentioned previously, this project is in the initial stages, and most of the activity so far has been planning and design. We have taken some protocols of searching print and current online encyclopedias to investigate strategies people employ, and have constructed some sample simulations of the sort that will be included in the encyclopedia (though none as grandiose as that presented in Section 1). The initial prototype system is being written in Lisp, running on Symbolics LM-2 Lisp machines (soon to be 3600's) and Xerox D machines (1100s), using pieces of representation code from RLL [11] and Eurisko [5,6], and simulation facilities from ThingLab [12].

From our viewpoint as AI researchers, the project offers a unique testbed for context-constrained understanding of speech, natural language, and other modalities which have not been programmable readily before (such as facial expressions and hand movements). It is a world which is almost as deep as most expert systems, and which is 500 times broader. Analogy and metaphor play a key role in human cognition, especially learning, but it has been difficult to experiment with those phenomena due to the narrow testbeds we have had available previously. We will exploit analogy as to facilitate data entry, and to explain exhibits to the user; other uses may be uncovered.

There is one final, crucial point to make. In the Abstract, we mentioned the Knoesphere's *eventual* impact on education, long after its introduction as a form of home entertainment. This delay is necessary because *browsing through a "science museum" is not an effective way to learn science.* As enjoyable as it might be, the system *presents facts,* it doesn't *teach principles.* After it exists, work can begin to add (probably manually) the effective teaching programmes, the critical analogies. Eventually, we hope that simulation-building tools will be good enough that, say, a physicist can expect to enter a simulation such as the Jovian example. This might well be done within the framework of guides and filters, but it is totally absent from current encyclopedias, and it will be largely absent from our initial system.

ACKNOWLEDGEMENTS

We thank Alan Kay - and Atari -- for setting up an environment that provides the resources and freedom to undertake a project of this size. We also thank Charles van Doren and Bob Stein.

REFERENCES

[1J Preece, Warren P. (ed.), *Encyclopaedia Britannica (15th edition),* Chicago, 1979.

[2] llayes-Roth, Frederick, Donald Waterman, and Douglas Lenat, (cds.,) *Building Expert Systems,* Addison-Wesley, 1983.

[3] e.g., NEX1S, the online text retrieval system by Mead Data Central.

[4] Sutherland, L, *Sketchpad: A man-machine graphical communication system,* MIT thesis, 1963.

[5] Lenat, Douglas B., "The Nature of Heuristics", *A.I. Journal,* Oct., 1982.

[6] Lenat, Douglas B., "The Nature of Heuristics III", *A.I. Journal,* March, 1983.

[7] Hayes, Patrick J., "The Naive Physics Manifesto", in (D. Michie, ed.) *Expert Systems in the Micro-electronic Age,* Edinburgh University Press, 1979.

[8] deKleer, Johan, and John Seely Brown, "Foundations of Envisioning", *Proceedings of the Second AAAI Conference,* CMU, 1982, pp. 434-7.

[9] McDonald, David, "Salience as a simplifying metaphor for natural language generation", *Proc. Second AAAI Conference,* CMU, 1982, pp. 75-8.

[10] Rich, Elaine., *Artificial Intelligence,* McGraw-Hill, 1983; see references to her work in this text

[11] Greiner, Russel, and Douglas B. Lenat, "A Representation Language Language", *Proc. of the First National AAAI Conference,* Stanford, 1980, pp. 165-9.

[12] Borning, A. H., "The Programming Language Aspects of ThingLab, a constraint-oriented simulation laboratory," *ACM Trans, on Prog. Languages and Systems, 3,* 4, Oct, 1981, 353-87.