

The Next Knowledge Medium

Mark Stefik

*Intelligent Systems Laboratory, Xerox Palo Alto Research Center, 3333 Coyote Hill Road,
Palo Alto, California 94304*

We are victims of one common superstition—the superstition that we understand the changes that are daily taking place in the world because we read about them and know what they are.

—Mark Twain
(from *About All Kinds of Ships*, 1892)

Warning: Flame follows.

—Gregor Kiczales,
(from an electronic network message,
August 1985)

The anthropological stories and the concept of memes were brought to my attention several years ago by Lynn Conway. Much of the vision and some of the material was drawn from a paper that we worked on together but never published. The important distinction between process and product, was made crisp for me by John Seely Brown, who also has encouraged and made possible projects like Trillium, which I watched with interest, and like Colab, in which I participated. Joshua Lederberg kindled my interest in biological issues and a respect for knowledge processes and their partial automation that has not faded. Dan Bobrow listened to my ramblings on several runs, agonized over my confusions, helped to get the kinks out of the arguments, and suggested the title for the article. Sanjay Mittal and I have spent many hours speculating together on the issues in building community knowledge bases and knowledge servers and in understanding the principles of knowledge competitions. Austin Henderson helped me to understand the Trillium story and to report it accurately. Austin and Sanjay hounded me to say, more precisely, what a knowledge medium is. Agustin Araya and Mark Miller participated in a Colab session in which we tried to jointly lay out these ideas, and together asked me to make the prescriptions clearer. Ed Feigenbaum persuaded me to be more precise in the discussion of the limits of today's expert systems technology.

Thanks to Agustin Araya, Dan Bobrow, John Seely Brown, Lynn Conway, Bob Engelmores, Ed Feigenbaum, Felix Frayman, Gregg Foster, Austin Henderson, Ken Kahn, Mark Miller, Sanjay Mittal, Julian Orr, Allen Sears, Lucy Suchman, and Paul Wallich for reading early drafts of this paper and for helping to clarify the ideas and improve the article's readability. Stephen Cross triggered the writing of this article when he invited me to give the keynote address at the Aerospace Applications of Artificial Intelligence Conference in Dayton, Ohio, in September 1985.

Public opinion about artificial intelligence is schizophrenic. "It will never work" versus "It might cost me my job!" This dichotomy of attitudes reflects a collective confusion about AI. What is AI anyway? How can we think concretely about what it is, what it could be, or what it should be?

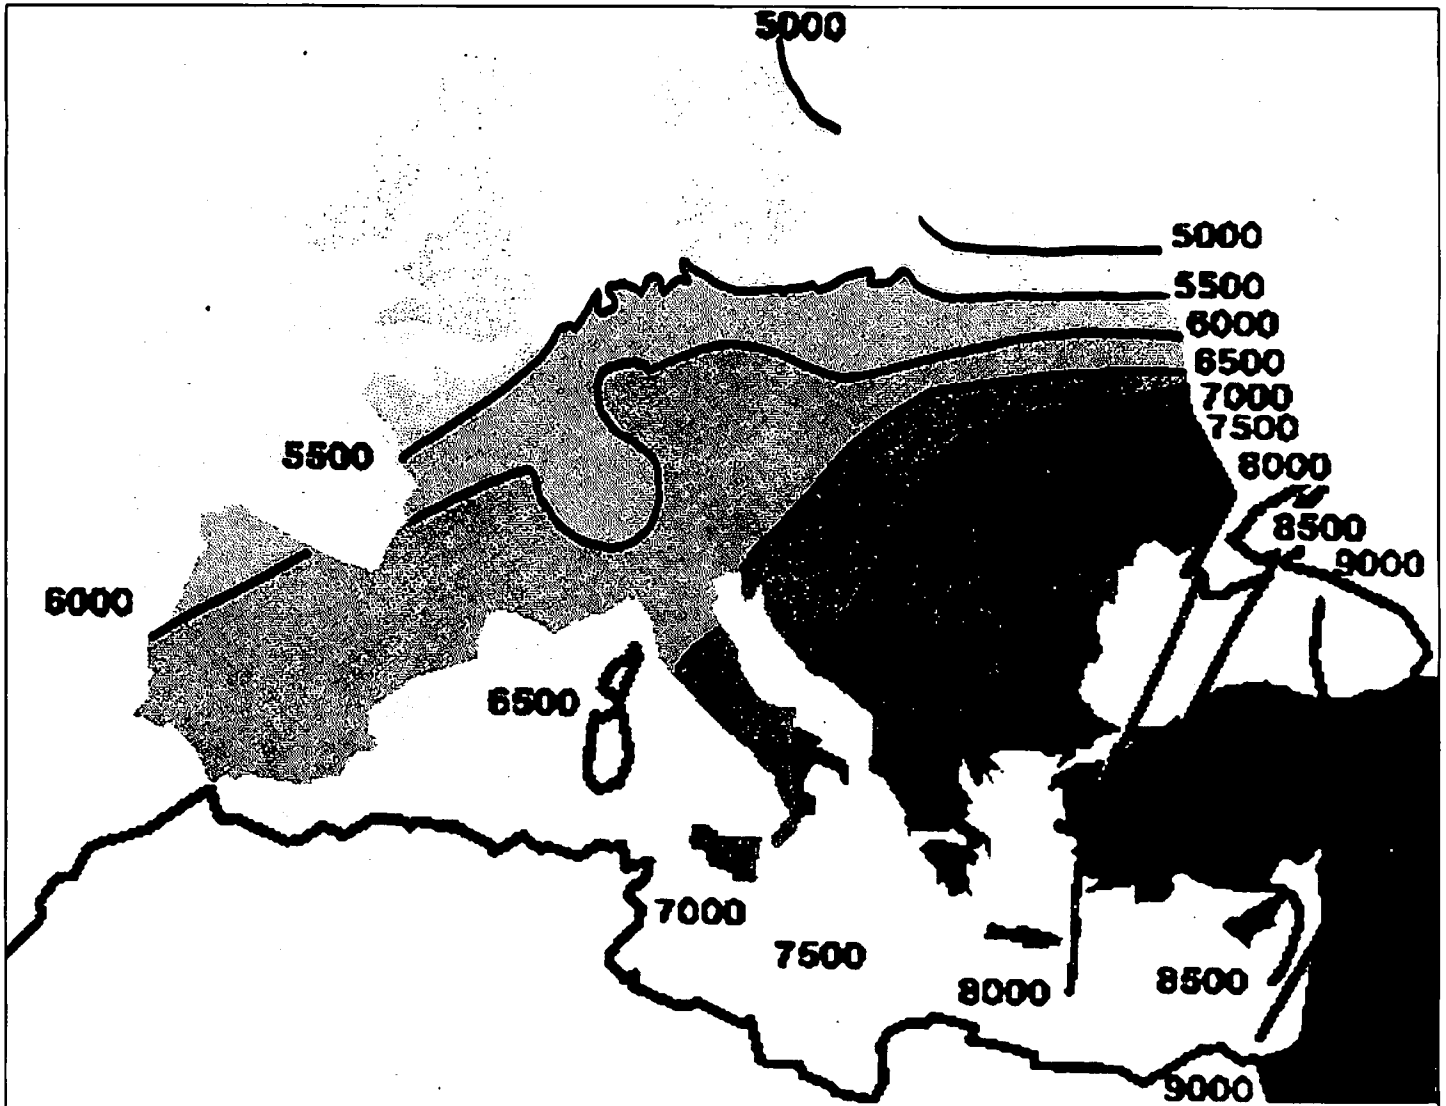
Most technologists are consumed with the activity of designing, building, and fixing things that need to work this year, if not next week. There is not much time for planning very far ahead. Nonetheless, futurists believe that AI will fundamentally change our way of life. Predicting the future is always a difficult and notoriously unreliable process, at least in specifics, but it is important to try to understand trends and possibilities.

This article examines how AI technology could change civilization dramatically. The article is in three parts: stories, models, and predictions. The stories describe processes of cultural change that have been studied by historians and anthropologists. They provide a historical context for considering present and future cultural changes. To illuminate these stories and their lessons about technology, several models of systems drawn from the sciences are considered. The models provide analogies and metaphors for making predictions.

Predictions are then made, drawing on some projects and ideas that might point the way to building a new knowledge medium—an information network with semi-

Abstract

The most widely understood goal of artificial intelligence is to understand and build autonomous, intelligent, thinking machines. A perhaps larger opportunity and complementary goal is to understand and build an interactive knowledge medium.



The spread of early farming from the Near East to Europe. The lines (called isochrons) indicate equal times of arrival of farming artifacts in the regions, as determined using radiocarbon dating. The artifacts reach Great Britain about 4000 years after they leave Egypt, indicating a cultural diffusion rate of substantially less than 100 miles for each century. (Figure reprinted with permission from A. J. Ammerman and L. L. Cavalli-Sforza, *The Neolithic Transition and the Genetics of Populations in Europe*. Princeton University Press, 1984, p. 59.)

Farming Culture.

Figure 1.

automated services for the generation, distribution, and consumption of knowledge. Such a knowledge medium could quite directly change our lives and, incidentally, change the shape of the field of AI as a scientific and engineering enterprise.

Stories

Mankind's cultures are constantly evolving. In the following three stories, the growth of knowledge and cultural change are considered. The stories are representative; literature contains many similar ones. They form a natural progression.

The Spread of Hunting Culture

At the end of the Pleistocene glaciations, a spear-throwing hunting culture swept from what is now the northwestern United States, throughout the length and breadth of North and South America. Paleo-Indian culture was characterized by the use of a spear with a distinctive fluted point and by the hunting of very large animals, such as bison and mammoths.

According to the archeological evidence, these artifacts, and presumably the culture, spread at a rate greater than one thousand miles per century. There is a debate as to whether the spread occurred through migration or

through cultural diffusion by the observation, imitation, and integration of the spear technology and hunting methods by tribes at the edge of a spreading cultural wavefront (isochron). The weight of the sparse evidence, however, favors migration of hunting bands following as-yet-unhunted herds. In either case, it illustrates the very rapid spreading of a prehistoric culture over long distances.

The Spread of Farming Culture

The second story involves the diffusion of early farming culture across Europe from Eurasia. The spread of farming culture has been mapped using radiocarbon dating on the oldest discovered farming artifacts in the regions. Figure 1 shows a map of the land area of Europe, Asia, and Africa, stretching from Germany and England southeast to the north shore of the Mediterranean and then south to Egypt. The first artifacts appeared in Egypt approximately 9000 years ago. Wavefronts tracing the progress of cultural diffusion are spaced at 500-year intervals across Turkey, Italy, Spain, Germany, and France. Farming artifacts reach Great Britain about 4000 years after they first appeared in Egypt. This diffusion rate is substantially less than 100 miles for each century (Ammerman and Cavallisforza, 1984).

It is easy to understand why the diffusion rate is so much slower for the farming culture than for the hunting culture. The farming culture is a much more complex form, containing several systems of knowledge. This culture brought about some dramatic shifts from food gathering and purposeful food collecting to the more organized behaviors of planting and harvesting. The collective investment in organization led to a great increase in population densities.

Peasants into Frenchmen

In contrast with the previous stories, this one illustrates that a technology can accelerate cultural change. This story is about the introduction of roads and railroads into France between the years 1870 and 1914 and the subsequent sweeping changes and modernization that took place. Similar stories of rapid change just prior to the twentieth century can be told for countries around the world, including the United States. The story of the modernization of France is well documented and easily available for those with an interest in history and technology (Weber, 1976).

Until the 1860s the highway system of France was a mere skeleton. Highways led to and from Paris, the seat of the central government; they were for troops to march on and allowed tax revenues to reach the treasury. The railway lines, which were begun in the 1840s, had the same characterization. They did not connect with the farms and villages and did not serve the needs of ordinary people. Most of the real traffic was on the trails, paths, tracks,

and lanes that covered the countryside. Along these trails traveled the people, goods, and ideas of the time.

A peasant's world was restricted to narrow corridors—the space of a village and familiar trails. Travel beyond the limits of a good hike was a difficult and costly undertaking. According to custom, the few who went to Paris, even if only once in a lifetime, were known as Parisians.

During the winter, the roads were so bad that they were classified according to how deep a person or a horse would sink in the mud—to the knees, to the shoulders, or to the head. Carts were unusable.


The exchange of goods was limited to neighboring regions. If a peasant wanted to sell wares at a distant location, that peasant faced the prospect of carrying them for hours and then arriving at a town with no means of storage. Once arriving, the peasant merchant was at the mercy of the buyers, who knew that the merchant was in no position to haggle. Consequently, farmers and regions tended to be self-sufficient.

The change came after 1881 when, in the public interest, a law was passed to promote the building of rural roads. Together, the railroads and the interconnecting secondary roads brought a new life to the villages, connecting isolated patches of countryside to the larger markets. The roads connected the villages together, and the railroads connected the nation together.

As the horizons of the peasant were expanded, new opportunities were perceived, and traditional orientations were abandoned. The peasant's apparent disinterest in trading evaporated. The necessary skills for shipping and receiving goods—reading, writing, and counting—were learned. The skills developed as part of a general education took on significance as occasions arose in which to use them.

Productivity expanded enormously. A rule of thumb of the time was that economic activity grew tenfold in any area serviced by the railroad. In the Correze region of south central France, for example, transportation made fertilizer available. Its consumption increased 13 times over, and crop production increased 65 times between 1866 and 1906. Industries were transformed as France began to function as a unified marketplace. The oldest of the cottage industries—spinning and weaving—was replaced by the new textile mills. Coke, transported by trains, replaced charcoal. The local nailmakers came into competition with nail-making machines. Transportation expanded the marketplace and made possible what we now routinely call the economies of scale. (See Figure 2.)

The railroads and roads of France, like those of other nations, triggered a process of rapid transformation. In a scant 40 years, they brought the French people a common market, a common language, a unified nation, and a new prosperity. In the words of many French politicians, roads were the cement of national unity. In Weber's words, the roads transformed "peasants into Frenchmen."



Technology can accelerate processes of cultural change. Until the 1860s the highway system of France was a mere skeleton. Travel was difficult and a costly undertaking, and trade was limited mostly to neighboring villages. In 1881 a law was passed to promote the building of rural roads, connecting small villages to each other and to the railroads. This triggered a process of invigorated economic activity, in which cottage industries were displaced by large-scale production in an expanded market. The roads also transformed France into a marketplace for ideas. The ideas crisscrossed France along the roads and railroads, bringing about new ways of thinking and a more uniform language. (Reproduced from the cover of *Peasants into Frenchmen* by Eugen Weber. Stanford University Press, 1976.)

Woodcut Print of the French Countryside.

Figure 2.

Models

The preceding stories illustrate processes of cultural change and the spread of knowledge. These processes can be accelerated by technology, as in the French roads example. Many other examples of technology accelerating change have been studied, such as the printing press, the post office, and the telephone.

In the following sections, four models of systems and change are examined: population genetics, ecology, economics, and processes in scientific communities. All of these models are well-known. They are reviewed here to

establish some terminology and metaphors that will be useful for making predictions about a knowledge medium.

Population Genetics

Beginning biology courses discuss the genetics of individuals, starting from the early plant experiments of Mendel to the more recent research that has revealed the chromosomal mechanisms of inheritance. Population genetics goes beyond the genes of individuals to consider the variations and percentages of genes in a population. The set of genes in a given population is called its *gene pool*.

In this model, living organisms are *gene carriers*, and a *species* is a single type of organism capable of interbreeding. Mutation occurs in the genes of individuals, and, over time, selection determines the percentages of different genes in a population.

As an environment changes, the selection processes also change; this is reflected by changes in the gene pool. *Genetic drift* refers to a change in the relative distribution of genes in a population. A fundamental hypothesis of population genetics is that when two groups become isolated from each other, there is always genetic drift between the two gene pools. Population geneticists have studied the mathematics of genetic drift and have related it to various factors such as the size of the gene pools. In general, larger populations have more stable distributions in their gene pools than smaller populations.

Sometimes, new species appear and displace related species much more rapidly than would be predicted by the apparent change of environment or expected rate of genetic drift. This phenomenon of speciation and displacement is called a *punctuated equilibrium* because of the characterization of a population that is stable for a long period of time and then experiences sudden changes. For example, the fossil record might show a sudden change in the shapes of teeth in a predator population.

The leading model for explaining this process takes place in three stages: isolation, drift, and displacement. First, a group becomes geographically isolated from the main population. In isolation, it undergoes selection and genetic drift across multiple genes more rapidly than does the larger body. Finally, the geographic isolation is removed, and the slightly fitter group competes against and displaces the original population.

The mathematics and concepts of genetic drift and gene pools can be adapted to other systems, even nonbiological systems. The systems must have replicating elements and analogous mechanisms for transmitting and recombining these elements.

Ecological Models

Ecology is the study of systems of organisms in an environment. The first observation from ecology is that systems have levels. This concept is perhaps best exemplified by the so-called *food chains*, in which big animals eat little animals and so on, down to the most rudimentary plants and microorganisms.

Levels are the most simple order of relations; in complex ecological systems, the relations between species form an intricate web. To describe these relations, ecologists have developed a rich vocabulary of terms: predators, symbiotes, parasites, and pollinators.

From ecology comes a familiar metaphor: the *ecological niche*. An ecology has many nooks and crannies for all of the functions that must be performed, and everything has its niche.

When several species evolve together in ways that increase their mutual adaptation, they are said to *coevolve*. From this mutual adaptation can come increased efficiency, and this leads to an important ecological principle: life enables more life.

In an uninteresting sense, everything depends on everything. This degree of interdependency does not mean that ecologies are fragile. Ecologies are not constant, nor are they formed all at once. They develop under processes of coevolution. Multiple species compete for and create niches. As an ecosystem increases in complexity, it also becomes more redundant and thereby more robust.

Systems with populations of replicating elements can be described by metaphors based on either population genetics or ecology. Population genetics provides metaphors for drift, mutation, and selection. Ecology provides metaphors for describing relations between groups of elements and coevolution.

Economic Models


Economic systems are similar in many ways to ecological systems. Businesses form economic subsystems and depend on each other according to intricate relations. There are suppliers, distributors, and consumers. Subcontractors supply parts and services to multiple manufacturers. Corporations and goods are said to occupy *economic niches*. New products can drastically change the shape of a market by displacing older products from existing niches or by creating different niches.

Economics brings us several concepts not found in ecology, including price, supply, and demand. These concepts provide a quantitative basis for explaining action in the marketplace. A market is said to "seek its own level," according to the laws of supply and demand.

When there are many suppliers in a market, some are more efficient than others, and consumers benefit from lower prices. Effects like this can ripple through an economy, as when a part that is used in many different products is made more cheaply. When there are many consumers, producers often can achieve economies of scale by switching to large-scale manufacturing processes and mass production.

Thus, business enables more business. This can also be seen in the spread of business. The first businesses in a rural area are basic and relatively inefficient. As a locality develops a rich mixture of businesses, its economy becomes robust.

Some economic systems require more sophisticated models than the essentially laissez-faire ideas just described. For example, an economy can interact with a legal system, it can be regulated, or more subtle phenomena might need explanation. Nonetheless, the above concepts form a reasonable first-order model for many situations and are enough for our purposes.



1894—Bridge Leading to the Eiffel Tower, Paris, France.

Stereoscopic view showing several forms of transportation in France. The Eiffel Tower was built in 1889. This view, from about 1895, shows barges, a horse-drawn cart, and a hot-air balloon. By this time, railroads and secondary roads had invigorated the French economy and had triggered a rapid process of change that made France into a unified nation. (Photograph by B. L. Singley. The Keystone View Company, Meadville, Pennsylvania.)

French Transportation. Figure 3.

Social Processes in the Scientific Community

Anyone with even a casual familiarity with science has heard about the *scientific method*. It is a principled approach for creating and validating new knowledge. Studies of the actual conduct of science, however, reveal a social richness in the conduct of science that goes beyond the scientific method.

Scientists are knowledge workers having important relationships with each other. Such relationships are the peer review and the “invisible colleges” of colleagues who share and collaborate on results prior to their publication. Some scientists are known to their colleagues as innovators. Some are best known for integrating the results of others or for reliably pursuing details. Some are best at theory, and others excel at overcoming difficult challenges in experiments. Some scientists are good at explaining things and contribute most with their teaching.

Thus, scientists have many different roles as knowledge workers. Science enables more science, or, perhaps, knowledge enables more knowledge. Yesterday’s discoveries and unanswered questions drive today’s experiments and provide the backdrop against which they are carried out. This additive effect is particularly evident in practical knowledge about the techniques of experimentation.

For example, a series of experiments about genes and nutrients might yield a well-characterized culture of microorganisms. This culture might then be used for fine-grained studies of genetic exchange, and these results, in turn, enable experiments about the mechanisms that regulate gene expression.

Predictions

In the last chapter of his book, *The Selfish Gene*, Richard Dawkins suggests provocatively that ideas (he called them *memes*) are like genes and that societies have meme pools in just the same way as they have gene pools.

The central theme of Dawkins’s work in biology was a shift to a “gene’s eye view” for explaining the processes of evolution and selection. When mammals reproduce, they do not clone themselves, and their offspring are not identical to them. Genes are the (mostly) *invariant* units of inheritance. Taking this point of view goes a long way toward explaining many of the persisting conundrums of traits and behaviors that are linked genetically.

Almost an afterthought in his book, Dawkins’s memes have been taken up by many writers since they were first introduced. Memes are carried by people. They are the knowledge units that are transmitted in conversations and

that are contained in minds. Memes can be reinterpreted in new environments and expressed in new combinations, just like genes. Memes, like genes, often come in clusters that work together. Memes compete for their share in meme pools, just as genes compete in gene pools. If populations of memes become isolated, they undergo *memetic drift*, analogous to genetic drift.

Toward a Meme's Eye View

Stories of cultural change can be reinterpreted from a meme's eye view. Cultural change occurs along a wavefront, with the memes competing and spreading to new carriers. Basic human capabilities for communication and imitation modulate the rate at which the memes spread. Differences in the rate of propagation, such as hunting culture versus the slower farm culture, can be explained by assuming that many more memes need to be communicated for farming than for hunting. Apprenticeship programs in science can be seen as a social mechanism for communicating memes about the techniques and practices of science that aren't reported in publications.

The progression in our stories from hunting culture to farming culture to the modernization of France is of increasing cultural complexity. It is not, however, a sequence of decreasing speeds of propagation. Considering only the complexity of the cultural shift, it might be expected that the modernization of France took many centuries, if not millennia, but, in spite of the dramatic cultural changes that took place, the rate of propagation sped up enormously. To understand this, the effects of the roads in France must be reconsidered.

The roads did more than change France into a marketplace for goods. They also transformed it into a marketplace for memes. The isochron waves, which so faithfully described the orderly flow of memes for the hunting and farming cultures, are completely inadequate for tracing the flow along the roads and railroads of France. Technology changed the process. Imagine the memes crisscrossing France, along the roads and railroads, creating an intricate pattern of superimposed cultural wavefronts. Ideas from faraway places were continuously reinterpreted and reapplied.

By bringing previously separate memes into competition, the roads triggered a shift in equilibrium. The relaxation of constraints on travel led to meme "displacement." Cottage industries were replaced by mass production, and the way of life changed. Multiple equilibria were punctuated at once. The very richness of this process accelerated the generation of new *recombinant memes* with their own wavefronts. Whole systems of memes (*e.g.*, how to run a railroad station, what the value of education is, and even how to speak French) were created and transmitted. As Weber noted, peasants became Frenchmen in a mere 40 years.

Since 1914 several new communication media have

been introduced, including improvements in the post office, telephones, and television. These communication media have quantifiable properties that govern the transmission of memes: transaction times, fan out, community sizes, bandwidth, and storage. Better post offices mean that people can spend less time traveling and that they have more time for other activities. Shopping can be done by mail. The rise of the mail-order catalog stores at the turn of the century is a manifestation of this change. Today in the United States, the large catalog stores connect consumers and suppliers into a large national marketplace that has tended to reduce regional differences.

AI Technology: Not Yet a Knowledge Medium

Precisely defining a knowledge medium is much like defining life, and, like life, it is better characterized in terms of processes rather than properties. Life usually is described in terms of processes such as reproduction, adaptation, growth, and the consumption of food. A knowledge medium is characterized in terms of knowledge processes such as the generation, distribution, and application of knowledge and, secondarily, in terms of specialized services such as consultation and knowledge integration.

For life there are many borderline cases that defy simple definition. Fires spread, change their burning patterns, increase in size, and consume fuel, but they are not considered living. Viruses and plasmids are classified as living because they take over the machinery of their hosts, lacking the machinery for reproduction. It can be said that mammals are "more alive" than viruses because the quality of their processes is so much richer. Knowledge media also have borderline cases: communication media without knowledge services and databases with limited distribution and services. Just as life is thought to have come from things that were "nearly alive," so too might knowledge media emerge from nearby media.

AI research includes topics relevant to knowledge media: the representation of knowledge in symbolic structures, the creation of knowledge bases for storing and retrieving knowledge, the development of problem-solving methods that can use and be advised by knowledge, and the creation of knowledge systems (or expert systems) that apply knowledge to solve problems.

However, AI technology, as it now exists, does not function in an important way as a knowledge medium in our society. Its influence has been far less important to the creation and propagation of knowledge than the secondary roads in France.

This is more than a matter of the youth of the field. The main goal of AI seems to lead off in a different, possibly contrary, direction. The term "artificial intelligence" expresses the most commonly understood goal of the field: to build intelligent, autonomous, thinking machines. Building autonomous thinking machines emphasizes ideas quite different from building a knowledge

medium. It suggests the creation of machines that are intelligent and independent. In contrast, the goal of building a knowledge medium draws attention to the main source and store of knowledge in the world today: people.

Books Versus Expert Systems

In the meme model, a *carrier* is an agent that can remember a meme and communicate it to another agent. People are meme carriers and so are books. However, there is an important difference: People can apply knowledge, whereas books only store it. Librarians, authors, publishers, and readers are the active elements in the printed knowledge medium. Computers can apply knowledge as well, and this makes them important for creating an active knowledge medium. When a medium includes computer systems, some of the knowledge services can be automated.

Knowledge engineers are the computer-literate monks of the twentieth century, illuminating their manuscripts in splendid isolation, awaiting perhaps the invention of the next printing press.

The most promising automated knowledge processors today are expert systems. They are the darlings of many high technology watchers; they are the business of several exciting start-up companies. In several well-publicized cases expert systems have proven to be of substantial economic value, far exceeding the cost of their development.

The tools for building expert systems continue to improve and the research tools of several years ago have become the programming and knowledge engineering power tools of today. We have our new AI Lisp machines or Prolog machines, and we have our knowledge base tools. Practicing knowledge engineers will correctly claim that these tools make a big difference. There are many anecdotal accounts of noncomputer specialists ("domain experts") successfully using these tools to build expert systems.

However, building an expert system is quite different from writing a book. In writing a book, an author needs to get the ideas together and to write them down clearly. Sometimes ideas will be missing or out of order or slightly wrong. But authors depend on the intelligence and knowledge of their readers to understand and integrate what they read. Not so with today's computers and expert systems. Today's computers are less sophisticated than humans. Knowledge must be acquired, represented, and integrated when programming an expert system. Moreover, the underlying tools, while providing assistance in the construction of the expert system (just as text editors provide assistance to an author) provide no memes of their own to help with the organization of new knowledge or to fill in its gaps. Each expert system requires careful hand-

crafting of its knowledge base, and for this reason expert systems are expensive.

An operational economy enables the process of manufacturing complex artifacts such as automobiles and airplanes. To succeed at making modern airplanes, a manufacturer exploits a marketplace for materials and sub-assemblies. A manufacturer does not need to also make tires and batteries, to mine metals, and to produce glass and plastics. A manufacturer does not want to master all the details of the necessary technologies; but it does want to exploit the economies of scale of the marketplace. Specialized companies can produce batteries, glass, and tires less expensively than can an automobile manufacturer. The marketplace makes it possible to build complex goods that would otherwise be infeasible if everything had to be done from scratch. This recalls the abundance rules from our models: life enables more life, business enables more business, knowledge enables more knowledge.

For complex systems such as airplanes and automobiles, the feasibility of manufacturing turns critically on the availability of high-quality low-cost goods in the marketplace. The "goods" of a knowledge market are elements of knowledge, or memes if you will. In today's expert systems, knowledge bases are built from scratch.

To return to our comparison of books and expert systems, both are highly creative enterprises, both require research to collect the facts, and in both cases there is no (or very little) economy of scale in writing N books or building N expert systems. Compared with the number of people who are literate in the printed medium, knowledge engineers are few in number. They are the computer-literate monks of the twentieth century, illuminating their manuscripts in splendid isolation, awaiting perhaps the invention of the next printing press.

Standardization and Shells

To reduce the cost of building expert systems, we need to be able to build them using knowledge acquired from a marketplace. This requires setting some processes in place and making some technical advances.

The technical issues are not just the usual problems of electronic connection; there are already networks for computers. Computer networks are used for many important tasks, such as booking airline reservations and clearing bank transactions. The networks carry mostly data, not knowledge; low-level facts, not high-level memes. This distinction eludes precise definition, but the general sense is that very little of what the computers are transmitting is akin to what people talk about in serious conversation.

Imagine drawing on a collection of knowledge bases for building expert systems. These knowledge bases would be developed for different purposes but would have some important terms in common. For example, consider the term "water." A chemistry knowledge base would specify when water freezes and boils and what dissolves in it. A

cooking knowledge base would include information about its measurement or its use with different kinds of utensils. A desert-survival knowledge base would relate water to sweat and night travel. Farming and boating knowledge bases would relay other unique information.

Anyone who has tried to give a computer program common sense has found that there is a staggering amount of it which is acquired on the way to becoming adults, and none of it is readily accessible to computers. Some AI researchers have started to build generic knowledge bases of potentially wide value. For example, there are projects in common sense reasoning and qualitative and naive physics. Lenat's CYC project at Microelectronics and Computer Technology Corporation (Lenat, 1986) is encoding the knowledge of an encyclopedia in an explicit knowledge base. The success of this project will depend on whether the entry of additional knowledge becomes simpler as the knowledge base increases in size.

AI has very little experience with combining knowledge from different sources. There is plenty of experience in using protocols for getting low-level information (bits) from place to place, and there is an established practice for encapsulating higher-level information above the low-level protocols. Experience with protocols does not go high enough, however.

A partial approach to combining knowledge from different sources is *standardization*. The goal of standardization is to make interchange possible. Initially, railroads were designed with different-sized gauges for different sets of tracks. By now, though, the diversity of railroad gauges has mostly disappeared, and, for the most part, railway cars can be routed along any set of tracks.

The idea is to create standard vocabularies (for example, using words like water) and ways of defining things in terms of primitives. This is the conventional approach used to build knowledge bases, where the transmission language is a simple transformation of the representation language.

Work on standardization can be coupled naturally to work on expert system shells. A *shell* is an environment designed to support applications of a similar nature. Shells are an intermediate point between specific applications and general-purpose knowledge engineering environments. Shells could be built for broad applications, such as planning, scheduling, and a variety of specialized office tasks. Shells have four things that knowledge engineering tools don't: (1) pre-packaged representations for important concepts, (2) inference and representation tools tuned for efficient and perspicuous use in the application, (3) specialized user interfaces, and (4) generic knowledge for the application. For example, a shell for a planning application would have representations that integrate multiple alternatives and beliefs with time. It would have interfaces for dealing with plans and alternatives. It would have generic categories for items such as time, tasks, and se-

rially reusable resources. It would include some generic domain knowledge, such as the fact that an agent can be at only one place at a time. Shells provide the potential for sharing and standardizing knowledge in communities larger than single expert system projects.

From Standardization to Interoperability

Shells and experiments with standardization are the right next steps, but they are only a beginning. Indeed, if standardization is the only approach taken for knowledge combination, it ultimately would defeat the whole enterprise. The fundamental problem is that memes are additive only when sufficient intelligence is applied for their integration. The fact that intelligence is needed to make knowledge "additive" is a lesson which has been painfully rediscovered several times. Early visions of the relevance of theorem proving to AI reflected this misconception. Much of the great appeal of building an artificial intelligence based on theorem proving was the notion that given a fast enough mechanical theorem prover, one could always add a few more facts and derive the consequences.

One lesson often cited from experience with theorem provers is that it is necessary for efficiency to be able to control search processes. There is an even more important lesson, however: A theorem prover is a fundamentally ignorant system. A child is often unable to make use of what it is told. Even an adult is often unable to work with ideas that are too far removed from familiar experience. Today's theorem provers know profoundly less than a young child; it is not realistic then to expect such systems to be able to integrate facts.

Nonetheless, this flawed notion has arisen in many visions of building intelligent systems. For example, a story similar to that about theorem proving could be told about the appeal and ultimate disappointment in schemes for encoding knowledge in terms of production rules. Additivity of knowledge requires more than a simple interpreter; it requires an intelligent agent.

When people read books, they actively integrate what they read with what they know. The process of learning from a book does not bear much resemblance to copying text (the "transmission language") to the end of a file.

A naive approach to developing standards for knowledge transmission is to repeat this mistake again, that is, to develop standard terms with simple fixed interpreters and to expect that, somehow, knowledge expressed in the transmission language will be additive. When people from different backgrounds share what they know, they must spend time mutually constructing a common background. Although standardization plays a role in shortening this process (because there is a substantial corpus of shared knowledge), people have also developed intricate means for discovering differences and developing communication.

Humans as knowledge carriers have developed techniques for "interoperability." The field of AI does not have

much experience in understanding how to do this, and this is an open area for research. Building shells and combining large knowledge bases might trigger such research by providing a setting and examples for exploring limits. The use of intelligent agents to compile and integrate knowledge into usable knowledge bases will highlight the difference between languages for knowledge transmission and knowledge representation. Transmission languages need to be rich in the descriptions that enable integrators to connect memes together; representation languages need to provide the necessary hooks for efficient access, introspection, and application of knowledge.


Roles in a Knowledge Market

Suppose there was a knowledge market. What would be the different roles in this market? From the economics model, we would expect to find knowledge producers, knowledge distributors (publishers), and knowledge consumers. Drawing on our practical knowledge of working markets, we might predict many other roles. Perhaps, there would be knowledge advertisers and knowledge advertising agencies. If the market were regulated, there might be knowledge certifiers. Experts who have participated in the creation of expert systems commonly report that the process of articulating their knowledge for explicit representation in computers has, itself, yielded a better body of knowledge and a more complete understanding of what they know. Reflecting on this experience, Michie has proposed the creation of *knowledge refineries*, where such processes could be used routinely to "purify" crude knowledge.

The model of the scientific community yields a different cut on the differentiation of roles. Integrators would combine knowledge from different places. Translators would move information between subfields, converting the jargon as needed. Summarizers and teachers would also be needed.

Workstations designed for professional knowledge integrators would need things unusual in today's AI workstations. An integrator needs to have ready access to the important knowledge media used in human affairs, so the workstation should provide technical bridges. It should include a scanner, so that books and journals can be read from their paper medium. The automated character recognition of text would not need to be perfect, because the integrator could help interactively with the rough spots. The process for converting a page of a book to a text file, however, should be convenient and mostly automatic. Similarly, it should be easy to scan in audio recordings or items from a remote database. The workstation should provide software tools for reorganizing the information, to aid the integrator in the profession of combining memes.

This vision of a knowledge medium might seem very distant in the future. There are a lot of objects and processes to create. How can such a process be bootstrapped?



Trillium is a sort of computer-aided design tool in which the controls for copiers can be designed. Trillium was not conceived as an expert systems project. Rather, it was built to augment existing media. It provides a language for expressing design concepts in terms of buttons, lights, actions, and constraints. Trillium quickly evolved into the best medium for describing user interfaces that Xerox copier designers ever had. Its use spread to several design teams within Xerox. It became a major (if not the major) medium for exchanging knowledge about user interfaces. Trillium has now been used for most of the major entries in the next generation of copiers made by the corporation.

Trillium Map.
Figure 4.

Bootstrapping a Knowledge Medium

The goal of building a new knowledge medium is not to replace work on expert systems with something else or to replace existing communication media. The goal is to tie these two elements together into a greater whole. A knowledge medium based on AI technology is part of a continuum. Books and other passive media can simply store knowledge. At the other end of the spectrum are expert systems which can store and also apply knowledge. In between are a number of hybrid systems in which the knowledge processing is done mostly by people. There are many opportunities for establishing human-machine partnerships and for automating tasks incrementally.

The best example of a knowledge medium using AI technology is the Trillium project at Xerox involving Austin Henderson and others. (See Figure 4.) Trillium created a knowledge economy of the memes of interface design for copiers. Modern copiers are being built with

many more powerful functions than the early machines, and the design of user interfaces for them has become much more challenging. From one perspective, Trillium is a sort of computer-aided design tool in which the controls for a copier can be designed. It provides a language for expressing the actions of a copier interface in terms of buttons, lights, actions, and constraints. Initially, Trillium was intended to facilitate the rapid prototyping of interface designs, so that designers could quickly try them and study the cognitive factors. Trillium quickly evolved into the best medium for describing interfaces that these designers ever had. Its use spread to several design teams inside Xerox, and soon these teams wanted to exchange their design concepts. Software to help them combine design concepts was developed and aided the teams in managing versions as they evolved differently at various sites. Trillium became a *major* (if not the major) medium for exchanging knowledge about user interfaces. It has now been used for most of the major entries in the next generation of copiers made by the corporation.

Trillium was not conceived as an expert systems project. Rather, Trillium was built to augment existing media. The computer integrates successfully with other existing media—in this case, phones and memos. What do computers and AI technology bring to bear in the Trillium project? The main benefit of Trillium is expressing the memes. In Trillium the memes of interface design are tangible artifacts in a knowledge medium.

Another project at Xerox that focuses on human collaboration is the Colab project, an experiment in inventing the team computer. This project was started with the observation that people spend much of their time in meetings, leaving their computers behind in their offices. The Colab is a computer-augmented meeting room in which computers are used as an essential medium in the process of meetings. Meeting software is developed for collaboration in organizing ideas for a paper and for arguing the merits, assumptions, and evaluation criteria for competing proposals. Tools support meetings through different phases, such as brainstorming, linking, and evaluating. The Colab also includes special hardware for group use, such as a large, touch-sensitive electronic blackboard. (See Figure 5.)

In both the Trillium and Colab examples, computers bring special capabilities not available in competing, passive media. In both cases, the computer provides active knowledge services, ranging from simple storage and retrieval of information to processes that interact with and augment human social processes. The languages provided by the tools encourage an important degree of precision and explicitness for manipulating and experimenting with the knowledge in that form. Trillium provides a substantially better language and communication capability than designers had previously. The Colab offers computational support for organizing information, file storage for saving

the information between meetings, and coordination that allows more than one person to write at the same time into a shared memory.

Revising the Goals of AI

Building a knowledge medium has a set of overall goals quite distinct from those conventionally embraced by AI. The vision of AI, suggested by its name, is the understanding and building of an autonomous, artificial intelligence. While building an artificial intelligence is compatible with creating a knowledge medium, AI breakthroughs are not a prerequisite for building or experimenting with them. Intelligence can be added incrementally to a knowledge medium. The enterprise of building a knowledge medium shares much of the technology that has become important to AI.

Goldstein and Papert (1977) announced a shift of paradigm in AI from a power-oriented theory of intelligence to a knowledge-oriented theory of intelligence: The fundamental problem of understanding intelligence is not the identification of a few powerful techniques but rather the question of how to represent large amounts of knowledge in a fashion that permits their effective use and interaction. The bottleneck processes in building expert systems are recognized to be getting knowledge into expert systems (knowledge acquisition) and subsequently modifying it and updating it over time. Recognizing this puts the field of AI in a position to shift even closer to the foundations of knowledge: from a focus on mechanisms of intelligence to the role of knowledge in intelligent systems to the augmentation of knowledge processes in a medium.

Agencies that fund AI could play an important role in promoting or accelerating this shift. For example, over the past few years, there have been many workstation projects that aim to support specialized knowledge workers. There have been projects for building workstations for physicists, engineers, geneticists, doctors, and others. These projects could be conceived as isolated applications of AI. In this case, with few exceptions, the projects will be quite narrow and won't become particularly large or important for the proposed client community.

Alternatively, these projects could be conceived in terms of building experimental knowledge economies and knowledge media. Much of the same work needs to be done in either case. The difference is a change of emphasis. An expert system project is usually (and conventionally) conceptualized in terms of an isolated and independent widget that carries out certain tasks. The conventional goal of AI leads to projects for which the creators can say, "Look ma, no hands!"

The knowledge medium requires a change of goal focus from product to process and introduces new criteria for evaluating projects. Important questions need to be asked: As a project evolves, where will the knowledge come from? How will it be distributed? How is knowledge distributed



Colab is an experiment in inventing the team computer, in which a group of collaborators can jointly construct and organize ideas. This project was based on the observation that people spend much of their time in meetings, leaving their computers behind in their offices. The Colab is a computer-augmented meeting room, in which computers are used as an essential medium in the process of meetings. The Colab also includes special hardware for group use, such as a large, touch-sensitive electronic blackboard. Meeting software is developed for collaboration in organizing ideas for a paper and for arguing the merits, assumptions, and evaluation criteria for competing proposals. Tools support meetings through different phases, such as brainstorming, linking, and evaluating. If the Trillium project and the Colab project are harbingers of things to come, then new projects for knowledge media will be conceived in terms of the means that they create for collaboration. (Photograph by Brian Tramontana.)

**Colab.
Figure 5.**

now? What kinds of knowledge will be distributed, and what form will this knowledge take? Will it be of a grain size that encourages recombination and synergy? How will multiple experts interact?

A funding program intended to build a new knowledge medium would include many kinds of projects. It would include experiments with small expert systems, conceived around shared community knowledge bases. It would include projects dealing with expert system shells and knowledge-transmission languages. It would include the development of low-cost multimedia workstations for knowledge integrators and research on the processes of knowledge integration. It would include traditional work on AI, for incrementally automating knowledge-processing tasks. It would include experiments in creating viable knowledge markets, with mechanisms for distributing and renting knowledge.

The original proposal for the Japanese Fifth Generation Project described a number of roles for knowledge-processing systems: increasing their intelligence so they can better assist mankind, putting stored knowledge to practical use, learning, and associating data. In his keynote speech for the second ICOT conference, Hiroo Kinoshita of the Ministry of International Trade and Industry hailed the creation of an advanced information society:

... in it, different information systems will be linked into networks, and a variety of services will be offered. In addition, rather than individuals playing the passive role of merely receiving information, they will be able to obtain that information which they require, use it, and transmit it among themselves, in what is expected to be a society more closely reflecting human nature (Kinoshita, 1985).

Such a network would be a knowledge medium in the sense that this term has been used here. As a policymaker, Kinoshita cited several difficulties in bringing this network and society into existence, including the long time to write software, mechanisms for computer (and knowledge) security, interoperability, and better man-machine interfaces. His concerns are valid and focus on technological limitations which would affect the processes that need to operate in a knowledge medium.

However, the Fifth Generation Project does not have any projects for building experimental knowledge ecologies or knowledge markets. If the Trillium project and the Colab project are harbingers of things to come, then new projects need to be conceived in terms of the means that they create for collaboration.

Building a knowledge medium is a long-term goal, complementary to the goal of building artificially intelligent agents. Importantly, the vision of a knowledge medium might be the more useful guide to progress. Like the agent goal, it is for the long term. It stands on other work in the larger field of computer science, such as work on databases and network technology. It rests on the same core work of AI—on language understanding, knowledge representation, and problem-solving.

Creating a knowledge medium relates directly to the human condition and raises fundamental research issues for ultimately creating elements of widespread value, such as community knowledge bases and semiautomated knowledge markets. If it makes knowledge accessible, it might continuously trigger minor punctuations of knowledge equilibria, as memes cross the relatively impermeable boundaries of human specialization.

Concluding Remarks

The AI systems of today are akin to the isolated villages of France before roads were built. Goods were made using time-consuming hand labor. The villages stood by themselves; in their poverty, they were relatively self-sufficient. Dialects were divergent, and experience was accumulated locally. There was little interest in the neighbors. The roads and larger markets were yet to be conceived and invented.

In the late 1890s, Robert Louis Stevenson persuaded the tribal chiefs of Samoa to cut a road through the wilderness. When it was opened, Stevenson said:

Our road is not built to last a thousand years, yet in a sense it is. When a road is once built, it is a strange thing how it collects traffic, how every year as it goes on, more and more people are found to walk thereon, and others are raised up to repair and perpetuate it, and keep it alive (Stevenson, 1896).

Stevenson's observation strikes me as profound; it illustrates a method for starting ideas or objects that will

persist. It clarifies the idea that a successful knowledge medium cannot be just an autonomous widget, but instead it should be a medium for seeding knowledge processes.

References

- Aceves, J. J. G. L., & Poggio, A. A. (1985) Computer-based multi-media communications (special issue). *Computer* 18(10):10-11.
- Ammerman, A. J. & Cavalli-Sforza, L. L. (1984) *The neolithic transition and the genetics of populations in Europe*. Princeton, New Jersey: Princeton University Press.
- Bobrow, D. G. (1985) Daniel Bobrow on artificial intelligence and electronic publishing. *Computer Compacts*, 6(3):55-57.
- Boyd, R., & Richerson, P. J. (1985) *Culture and the evolutionary process*. Chicago: University of Chicago Press.
- Brown, J. S. (1983) Process versus product: a perspective on tools for communal and informal electronic learning. In S. Newman & E. Poor (Eds.), *Report from the learning lab: Education in the electronic age*. New York: WNET Educational Broadcasting Corporation.
- Crane, D. (1972) *Invisible colleges: Diffusion of knowledge in scientific communities*. Chicago: University of Chicago Press.
- Dawkins, R. (1976) *The selfish gene*. New York: Oxford Univ. Press.
- Dobzhansky, T., Ayala, F. J., Stebbins, G. L., & Valentine, J. W. (1977) *Evolution*. San Francisco: W. H. Freeman.
- Feigenbaum, E. A. & McCorduck, P. (1983) *The fifth generation: Artificial intelligence and Japan's computer challenge to the world*. Reading, Mass.: Addison-Wesley.
- Fleck, L. (1979) *Genesis and development of a scientific fact*. Chicago: University of Chicago Press.
- Friedland, P. (1985) Special section on architectures for knowledge-based systems. *Communications of the ACM* 28(9):902-903.
- Goldstein, I., & Papert, S. (1977) Artificial intelligence, language, and the study of knowledge. *Cognitive Science* 1(1):84-123.
- Hookway, C. (1984) *Minds, machines, and evolution: Philosophical studies*. Cambridge, England: Cambridge University Press.
- Kinoshita, H. (1985) Towards an advanced information society. *ICOT Journal*, 6:9-19.
- Lewin, R. (1985) Punctuated equilibrium is now old hat. *Science* 231:672-673.
- Lederberg, J. (1978) Digital communications and the conduct of science: The new literacy. *Proceedings of the IEEE*, 66(11):1313-1319.
- Lenat, D., Prakash, M., & Shepherd, M. (1986) CYC: Using common sense knowledge to overcome brittleness and knowledge acquisition bottlenecks. *AI Magazine*, 6(4):65-85.
- Michie, D. (1983) A prototype knowledge refinery. In J. E. Hayes & D. Michie (Eds.), *Intelligent systems: The unprecedented opportunity*. Chichester: Ellis Horwood.
- Moto-oka, T. (1981) Fifth generation computer systems: Proceedings of the international conference on fifth generation computer systems. Amsterdam: North-Holland.
- Price, D. D. (1983) Sealing wax and string: A philosophy of the experimenter's craft and its role in the genesis of high technology. Washington, D. C.: American Association for the Advancement of Science.
- Stefik, M., Foster, G., Bobrow, D. G., Kahn, K., Lanning, S., & Suchman, L. (1986) Beyond the chalkboard: Computer support for collaboration and problem solving in meetings. (In prep.)
- Wallich, P. (1985) Computers and the humanities. *The Yale Review*: 472-480.
- Weber, E. (1976) *Peasants into Frenchmen: The modernization of rural France, 1870-1914*. Stanford: Stanford Univ. Press.
- Wilson, E. O. (1980) *Sociobiology: The abridged edition*. Cambridge, Mass.: The Belknap Press.