

I Had a Dream: AAAI Presidential Address, 19 August 1985

Woody Bledsoe

Microelectronics and Computer Technology Corporation, 9430 Research Boulevard, Austin, Texas 78759

Twenty-five years ago I had a dream, a *daydream*, if you will. A dream shared with many of you. I dreamed of a special kind of computer, which had eyes and ears and arms and legs, in addition to its “brain.”

I did not dream that this new computer friend would be a means of making money for me or my employer or a help for my country—though I loved my country then and still do, and I have no objection to making money. I did not even dream of such a worthy cause as helping the poor and handicapped of the world using this marvelous new machine.

No, my dream was filled with the wild excitement of seeing a machine act like a human being, at least in many ways.

I wanted it to read printed characters on a page and handwritten script as well. I could see it, or a part of it, in a small camera that would fit on my glasses, with an attached earplug that would whisper into my ear the names of my friends and acquaintances as I met them on the street. Or in a telephone that allowed me to converse with a friend in Germany, he in German and me in English. For you see, my computer friend had the ability to recognize faces, synthesize voice, understand spoken sentences, and translate languages, and things like that.

I’ll admit that in 1960 my computer person had a much larger head than I envision for it now. Because I then didn’t know about microcomputers.

My dream computer person liked to walk and play Ping-Pong, especially with me. And I liked to teach it things—because it could learn dexterity skills as well as mental concepts. And much more.

When I awoke from this daydream, I found that we didn’t have these things, but we did have some remarkable computers, even then, so I decided then and there to quit

I want to thank Doug Lenat, Mary Shepherd, Clive Dawson, Joe Scullion, Hassan Ait-Kaci, Elaine Rich, Dick Martin, and Dick Hill for helpful comments. The title is obviously due to Martin Luther King

my job and set about spending the rest of my life helping bring this dream to reality.

Sometimes I forget this dream, and for these periods my life is more drab. Recently a reporter asked me, “Why do you scientists do AI research?” My answer, “Well certainly not for money, though I wouldn’t mind being rich. It goes deeper, to a *yearning* we have to make machines act in some fundamental ways like people.” Period. This reminded me of my dream—and brightened my life again for awhile.

But you know this is not a fairy tale. I *do* want to see *my* computer prove difficult new theorems in mathematics, recognize faces, diagnose diseases, reason by analogy, *learn* in many ways, engage in intelligent discourse, use common sense, play a good game of bridge and tennis, etc., etc., etc.

It is amazing how this seemed so possible in 1958, 1960, 1962. You know this still seems possible to me now; that’s what *really* drives me. (Of course, some of it has happened.)

There is a buck to be made in AI now. And I don’t mind that; in fact, the economic interest might just push forward part of the needed research. But I believe that it is not that buck (dollar, yen, franc, pound, mark, lira, peso, . . .) that drives the few of us who can and will *really* make it happen. We march to a different drummer, and I am proud to be part of that battalion which responds to *that* music.

I don’t have to explain why I had (and have) this dream; I just do. It does not require me to be a mathematician or an engineer or computer scientist. I just want the results; well, almost—what I want, too, is to be surprised. It does not even require that we do “good science,” though that would probably help if it is not overdone.

My dream is Buck Rogers, HAL, 2010, Star Wars, R2D2, CP30, the Turing test, Objects of the third kind, Jules Verne, all piled into one, but with all the ridiculous things divided out, like breathing in a vacuum and going faster than the speed of light, leaving what is somehow

possible by our present knowledge of science. That is what is exciting—doing the possible.

And I am not so generous; I want to actually see these things myself before I die. “Pressing science along a little bit” won’t cut it with me unless I am part of the excitement and see some of the major milestones.

The physical parts of the proposed computer person seemed important in 1960. And this still excites me now—to have a robot run, fall, and get up; to act autonomously, to be a moving companion. Oh, I’m well aware that the real problems are those of the mind, getting computer programs to act as if they reason, act as if they understand, think, learn, plan, enjoy, hate, etc. That is the challenge of the age. Some of this has happened already, and I believe that we have, in this room, the talent to bring about much of the rest of it. As it unfolds during the next years and decades, let us not fail to stop occasionally to enjoy it, to “smell the roses,” to thrill as each new milestone is reached.

When I began to prepare for this talk, I wondered what I would have to say that was worth hearing. And it was this contemplating that brought to mind the above thoughts, this “dream.” These are things that are important to me in my “calling” as an AI researcher.

These 25 years have not been totally kind to my dream: Shakey liked shaking more than running and thinking, and was laid aside for a season; language translation sputtered, died, and was resurrected; facial recognition was pushed back on the researcher’s stack; automatic provers showed signs of growing pains, which disheartened the fainthearted; no machine stepped forward to try the Turing test; robot arms were duplicating block castles instead of playing squash; etc., etc., etc.; many AI researchers lost faith and dropped out.

But, curiously, I remained in the AI fold, and why? Because these 25 years have also been fruitful and exciting. We have much to be proud of, with much left to be done.

First and foremost, we have learned *what* we are. Just as a small child remains ineffective until taught (given knowledge), so it is with our machines. Reasoning alone could not have enabled a prehistoric man to even invent the wheel, no matter how nimble his brain, but the space-age woman with her *knowledge* of wheels, gears, engines, computers, aerodynamics, and the like, with the same reasoning power, can discover much more. Because knowledge is *king*; knowledge—the key to who we are. Even reasoning itself is enhanced by knowledge about reasoning and knowledge about what we are reasoning about. Ours is, in essence, the knowledge business. (Ed Feigenbaum says that we are working on “Knowledge-application machines.”)

But you say, “Every ‘smart’ program, used to solve a problem, regulate a chemical process, design a bridge, etc., has key pieces of knowledge built into it. So what is new about AI?” The answer is that the AI scientist or engineer

recognizes this knowledge for what it is and has, in the case of expert systems, plucked it out of the program and placed it in a separate “knowledge base.” Not only does the knowledge give the power, but it provides the functionality. The knowledge base acts as a new and powerful computer language, which is used by the programmer to carry out his or her will. He/she defines functionality and causes actions merely by changing this knowledge base.

So, foremost, we have learned that we must use knowledge, the knowledge accumulated by mankind over these last few thousand years, if we are to achieve these AI dreams. And we have accomplished a great deal during these last 30 years; let me mention some of it.

First, let me express my annoyance with some of the distracted individuals who criticize AI researchers for not “jumping to infinity” in one leap. Somehow, to them it is OK to work step by step on the dream of obtaining controlled thermonuclear energy or a cure for cancer or a cure for the common cold, but no such step-by-step process is allowed for those trying to (partially) duplicate the intelligent behavior of human beings. To these cynics, a natural language system that converses with us in a restricted form of English is somehow not a legitimate step toward passing the Turing test. I know of no case in the history of science where such “naysayers” actually helped with a new discovery.

Indeed, almost all of our AI accomplishments have been of the *partial* kind: natural language processors that handle a *subset* of English (or French, etc.); systems that recognize and synthesize *limited* forms of speech; character recognition machines that read only *typewritten* characters; expert systems that perform a variety of tasks (but not all that a human can); theorem provers that can prove difficult theorems in a *particular* area of mathematics or that can handle the inferencing needed for elementary expert systems, including nonmonotonic reasoning; programs that play expert-level chess; programs that exhibit an elementary level of learning and reasoning by analogy. And the list goes on.

Another key thing that we have learned and are still learning is the list of crucial technologies needed to continue the pursuit of our AI objectives. These partial results, mentioned above, have helped to unearth the roadblocks that stand between where we are now and where we are trying to go with AI. We are beginning to enumerate and classify these enabling technologies.

Foremost on the list is the representation and storage of knowledge, with the added requirement that the particular design will allow:

- Learning: Ease of acquiring and storing the knowledge
- Performance: Effectiveness in using the stored knowledge to perform tasks, solve problems, and answer questions

I believe that it is time to build large, very large,

knowledge bases. Such a knowledge base should contain common sense knowledge as well as encyclopedic and expert knowledge and be structured to handle the learning and performance requirements mentioned above. (An effort of this sort headed by Doug Lenat at MCC uses common sense knowledge in a fundamental way and uses analogy to help with knowledge acquisition and problem solving.)

It is believed that such a large-structured knowledge base would not only allow the sharing of knowledge by numerous systems, but, if structured correctly, could provide much more robustness and functionality than is possible from a number of distinct smaller KBs.

It has been said that we cannot have true machine intelligence until we have *effective* machine learning. In that case, we have a ways to go. But a number of good researchers are beginning to make progress in this area. Earlier work on machine learning tended to be too ambitious, too general, whereas the recent efforts have had more success where the things being learned are controlled by knowledge structures, where the machine finds values of facets within a human-supplied framework. But even so, until we see some *real* gain, a reasonable amplification of capacity, then some of us need to be rethinking the learning and analogy process from scratch. (This rethinking might also apply for some other areas of AI research.)

Twenty years ago one might have been tempted to say that it requires only two things to build a machine which appears to think like a human being—*machine learning* and *natural language understanding*—because such a machine can be taught by feeding it more and more knowledge from existing books, letting it bootstrap itself to higher and higher levels of mental functionality. But those two requirements are formidable indeed. In fact, I'm afraid that this characterization is misleading. It lets us believe that the major present needs of AI can be satisfied through machine learning. While that might be a correct principle for the long run, it won't do for the near term. So we must press on in other areas of AI as well as machine learning. For example, the important work on speech acts should be pushed *now* and not delayed until the machine can learn from books how to carry on a discourse.

I like Marvin Minsky's suggestion that the ability of a program to learn should be proportional to what it already knows. Such a program, when and if it is achieved, can be exploited in a dramatic (frightening?) way.

Causality is another important research area in AI. As our intelligent programs such as expert systems begin to fail, we want to move from "shallow" (statistical) rules toward "reasoning from basic principles." Several research programs are pushing in this direction. I believe the key here is to move *toward* basic principles, a step at a time, and not *to* basic principles in one step. For example, knowledge of actions can be classified by levels of causality. I will try to explain this by first giving an example. If

one holds an object in his hand five feet above the ground and releases it, it will

1. Fall toward the ground
2. Fall toward the ground with increasing velocity
3. Fall, with its height y , in feet, governed by the equation

$$y = 5 - 16.1t^2$$

with t measured in seconds

4. Fall according to Newton's law of attraction
5. Do the same as (3) but also account for air friction
6. Fall according to the laws of general relativity

For most applications, the first answer is enough: "If you release it, it will fall." For example, we might say to a child, "If you drop that rock it will hurt your foot." This might be called the "shallow" level. Deeper levels give information that is more and more precise but at a higher cost.

A human could never get anything done operating continuously at the third level, let alone the fifth, and neither could an expert system. The early expert systems tended to operate at the first level of causality, the shallow level. This was fortunate because it allowed these programs to exhibit a great deal of expertise for minimum cost. Such successes of expert systems have been of great value to the field of AI. They not only helped build confidence in the AI researchers that worthwhile accomplishments are possible but also promised financial returns in the near term that can help pay for further research and development. So operating at the shallow level is not bad at all when it works.

The problem comes when that level is not adequate, when deeper causal reasoning is needed. And it is at this point that our machines need to be directed one step deeper in the causality chain.

Using causality properly, then, does not mean jumping to the deepest causal level, but rather working down through levels as needed. I believe that the recent work on qualitative reasoning is a correct step in this direction. But to make it work properly, an overall knowledge structure, governed partly by common sense, is needed to control the process.

These new super expert systems (for the coming decade) will absorb a large percentage of the research and development effort over the next several years, and rightfully so. I mean expert systems that have been endowed with large structured knowledge bases, ability to reason through various causality levels (preferring the shallowest but resorting to deeper levels as needed), limited ability to learn automatically from experience and to accumulate knowledge by analogy, truth maintenance systems,

enhanced human interfacing to facilitate knowledge acquisition from experts and for ease of use, etc.

These super expert systems will evolve into the “thinking” part (as opposed to the moving and sensory parts) of our dreamed of intelligent machines of the future. Later versions will have enhanced ability to learn (*e.g.*, learning directly from machine-readable text) and reason by analogy and much more.

Now let me list some of the areas that I feel will dominate AI research over the next decade. I have discussed most of these already and now list one more: automatic reasoning.

Important Research Areas

Large-Structured Knowledge Bases

Knowledge Representation

Knowledge Storage, Retrieval, and Use

Expert Systems Technology (A large effort)

Machine Learning

Controlled by knowledge structures

Causality

By depth levels

Human Interfacing

Natural Language Processing

Speech Recognition and Generation

Automatic Reasoning

Analogical Reasoning

Common Sense Reasoning, Default Reasoning

I have not tried to be complete in this listing and have not even mentioned some important areas such as robotics, automatic programming, and planning.

Automatic reasoning is an area of research that is becoming increasingly important for a number of reasons. Earlier expert systems required only modest inferencing power because they operated on rules at the shallowest levels. But as we reach toward deeper causality, the reasoning component is challenged to handle the switching of levels and the added complexity of the deeper levels. In this, as always, knowledge plays a crucial role.

Also, the emergence of logic as a basis for programming languages (PROLOG, LOGLISP, PARLOG, etc.), and as a means for storing knowledge (in logic databases, logic based rules for expert systems), has suddenly placed a new load on our automatic-reasoning programs (our provers). Thus we see the great interest in “Kilolip” machines that perform a large number of logical inferences per second. Such high performance will not only be needed for horn-clause problems, such as the use of PROLOG, but also for reasoning in first-order logic and even in modal logic and higher order logic. Thus, the renewed interest in automated theorem proving.

It will be interesting to see whether the new concepts for handling horn clauses and first-order logic, which are expected to produce “raw horsepower” in the Megalip range, will be enough to cope with the load that will be imposed by the next generation applications, or whether these methods will have to be “spiced” with special reasoning knowledge units for speeding up proofs for particular applications. In any event, automatic-reasoning research should become more relevant in the near future.

Let me not be misunderstood. General-purpose reasoning machines (theorem provers) alone are not enough. Knowledge is still the key. But the requirements for reasoning about knowledge will be intensified and partly satisfied by these new high-speed provers that are beginning to appear.

I have great faith that the AI community is headed generally in the right direction. About half of the new crop of graduate students admitted to the Ph.D. program in computer science at the University of Texas this year selected AI as their preferred field of study. This preference for AI seems to be duplicated throughout the world, and we are talking about some of the very best students. These young people hold in their hands the future of this discipline. The power and influence of the earlier pioneers will wane as these new researchers emerge.

I urge these new students and all new researchers to set themselves a vision of the future and to have the courage to make major new departures, to question the old and get on with the new. There is much to learn from us. We have pointed generally in the right direction, but the major gains are yet to be made.

I personally favor the bold approach over the timid. And there are certain bold experiments that have to be made. One such effort was the mechanical translation (MT) work of the early 1960s. Some have called it a failure, but I do not! It had to be tried. It seems rather obvious now that you cannot have MT without language understanding. That awareness was made much clearer by these earlier experiments—they helped focus research in the important area of natural language processing. And look how exciting that has become and where it has led, even to the resurgence of MT! A similar story could be told about early speech recognition—quality speech recognition is not possible without language understanding. Early experiments with perceptrons represent another such example. In all these cases a lot of work compensated somewhat for the lack of a great idea. (The Shakey robot project at SRI is another example, but in that case the value of the early work is widely appreciated.)

The principle I want to make is this: when you have what looks like a good idea, give it your best shot, waste a little money to get some early feedback. Don't take forever to study the problem, because that is even more expensive (and less exciting). Of course, this strategy (this scientific method) requires *character* on the part of the researcher.

He/she must be willing to analyze those experiments, reformulate theories, and press on. Otherwise, that person does not qualify for the work and should not be entrusted with research funds.

I was recently reading about Thomas Edison and his team at the time they developed a successful light bulb. He started with what he thought was a good idea and plowed ahead. He was brash, he was cocky, he bragged about what they would do (build a widely usable electric light bulb), his early ideas were wrong. I believe that they were lucky, even with all their brilliance; it could have taken years. But this is another example, like language translation, where an early expensive failure returned information that helped finalize a successful solution.

I believe that AI is in a position today where these kinds of bold experiments are needed (but not the bold bragging). They need to be conducted by men and women with character, with wisdom and persistence enough to succeed.

Another concern I have is the "flash in the pan" researcher, the person with a limited theory, who does a trivial application of it, or none at all, and gets no useful feedback. He builds a program that cannot surprise him in any way and leaves to others to prove and extend his work. His fragment had better be pretty brilliant if anything is to become of it. More likely a *real* researcher will rediscover the fragment as part of a larger effort and absorb most of the credit. We might recall that most AI pioneers are well known for what they *did*, not what they theorized.

What is the most important characteristic of a good researcher? Answer: He does good research. Successful people somehow find a way to succeed, others fail. Of course, native intelligence is an important ingredient, but it alone is not enough. An equally important characteristic is the ability (and inclination) to combine theory and experiment.

So, again, I would say to young people, set a dream. Set a goal (your part of bringing about that dream). Tool up: education, employment, facilities. Pursue it with vigor—and impatience. Want it today. I've never seen a content researcher who was worth his salary. Don't be easily deterred by those who don't have your insight and training. Work hard, provide momentum, don't give up easily. Don't spend too much time extolling the work of others; you will never be properly recognized or satisfied until you make your own personal contribution. Compare and compete. These are rules for a researcher in any field. My conviction is that the field of AI is worth your finest efforts.

I have told you about my dream, have offered advice for young researchers, and have offered my opinion on important areas of AI research. But of all the predictions that I could make, the one that I'm most sure about is that we will again be *surprised*.

**WHAT SERVES AS THE TWO PRIMARY
COMMUNICATIONS LINKS WITH THE
ARTIFICIAL INTELLIGENCE COMMUNITY?**



The **AI MAGAZINE** and **AAAI PROCEEDINGS**.

Learn more about the **AIMagazine**, **AAAI Proceedings**, **AI Conferences**, and other benefits associated with joining the **AMERICAN ASSOCIATION FOR ARTIFICIAL INTELLIGENCE (AAAI)** by calling or writing the **AAAI**, 445 Burgess Drive, Menlo Park, CA 94025, USA, (415) 328-3123



American Association for Artificial Intelligence

**Call for Workshop Program
Sponsored by
The American Association
for Artificial Intelligence**

The AAAI has supported small workshops for the last several years. This support includes publicity, printing, office help, and subsidies for other expenses. \$5,000.00 is a typical subsidy, but up to \$10,000.00 may be considered. Any topic in AI science or technology is appropriate. Anyone may volunteer to organize a workshop on any topic. The organizer(s) should determine the topic, the date, the site, and the procedure for selecting papers and attendees. S/he should also decide whether preprints should be distributed.

Proposals for scientific workshops should be made to:

Professor John McCarthy
Computer Science Department
Stanford, CA 94305
(415) 497-4430,
jmc@su-ai.arpa

For workshops on applied topics, applications should be made to:

Dr. Peter Hart
Syntelligence
P.O. Box J620
Sunnyvale, CA 94088
(408) 745-6666
hart@sri-ai

AAAI proposes that program committees give special consideration to papers that have been presented at workshops in choosing invited speakers for national conferences.