Issues & Opinions

The Mature Intelligent Computer

There is a good deal of speculation about the possible nature and capabilities of computer systems employing artificial intelligence—some of it pretty wild. The possibilities are many, but they are not unlimited. Limits can be inferred from studying the history of the field. Most aspects of AI are not, in fact, new; for instance, my first forecast of the prospects for speech recognition technology was prepared more than 20 years ago. And expert systems, which receive most of the attention now, date back in demonstration form to the early days. (I saw a computer learn the optimal strategy for playing the game "Nim" in 1959.) More limits can be inferred from observing the history of innovation in software, my observations of which span more than 30 years. I think it is possible to make sensible speculations about the future of artificial intelligence for the next two or three decades—personal opinion to be sure, but informed.

New hardware won't have much to do with it. For the next two decades intelligent computers will still rely on magnetic and semiconductor technology (with maybe some optical disks). Exploratory research is being performed on two new types of component technologies; optical and organic. Optical technology (in which light impulses are processed instead of electrons) may be able to run faster than semi-conductors, but it uses more energy. Circuits will have to be less dense and more complex. Organic technology (in which circuits are grown, not made) may be dense, but cheap and the idea of growing computers with components as small as nerve cells is fascinating. However organic components are much slower than electronic ones.

Much of current artificial intelligence research concentrates on machine recognition of human speech and of visual images, recognition of the **meaning** of sound or sight patterns. The difficulty with pattern recognition is separating signal from noise—identifying which of the incoming signals are related to one another and to the meaning being sought. Humans do this very well. We can concentrate on a single conversation in a crowded restaurant or recognize a face in a crowd. So far, however, we have been unable to tell machines how we do it.

Recognition of visual patterns of images on paper or in electronic signals (television, radar) presents a similar problem with explicitness. A *direct* pattern match can determine meaning. This may mean that robots can depend on artificial vision in factory workstations where the objects in the visual field can be limited. But a robot housemaid, able to take proper action on the basis of the unconstrained sights that may appear in a home, is as far away as ever.

In laboratories we can analyze speech sounds (spectrum, power density and the like) and compare the resulting patterns against pre-stored patterns to see what word is being spoken. This works poorly because too many speech sounds are the same. Machines can recognize only a small set of discrete commands or numbers, even when efforts are made to consider context. The brain does no such sound analysis and pattern matching. Instead, it somehow interprets the meaning of the incoming sounds continuously in terms of context, past associations and emotional keys. Twenty years of research has produced remarkably little progress in determining just how the human brain interprets speech. The prospects of machines able to extract meaning from continuous, unconstrained speech seems poor, even into the 21st century.

More limited forms of speech recognition would still be useful. For instance, a machine that could type properly spelled and punctuated text from oral dictation would be a handy gadget, even though ignorant of the meaning of the material. Such a "speech speller" may be possible based on an examination of the surrounding words and sentence structures. And machines that can understand meanings of carefully constrained vocabularies (such as those used by pilots, machine operators and artillery spotters) should also be useful.

Users may converse with future computers in a new form of command language, not in the spoken languages of humans. It will be new because it must match the always-explicit computer with the often-ambiguous human. No human language has had to cross this barrier. It is already clear that some of the expressions used in the language will be without human precedent (e.g.,

the "icons" used in commanding many personal computers). People will probably learn a generic version of this language while they are children in school. They will then be able to quickly adapt to the dialects of command language they need to use as adults, much as an experienced driver can adapt to an unfamiliar car.

Communicating with computers will become much easier. Computer programs will learn what the user wants by asking questions and accepting corrections, remembering the results, and being able to provide the same service the next time without repeating the instructions. To a degree they do this now—operating systems and communication control programs for most computers require set-up only once and then remember the setups. In many instances they provide prompts and ask questions to help the user remember what needs to be specified.

They will be doing a better job in the future. A few information retrieval systems are already being used that accept user interest profiles, match them periodically against new material, and automatically send the user material that matches the profile. Such a system is anticipating the user's needs based on all previous instructions. The user can keep refining and amplifying the instructions without limit.

Expert systems should help even more in selected areas by adding the expertise of other humans to the system's ability to anticipate user needs. In an expert system there is a "knowledge base" of facts, decision rules and diagnostic questions that has been built up by human experts. The user conducts a dialog with the system in which it helps him define the problem, then solve it if the knowledge base contains the solution (if it does not, it will tell the user what is lacking). The expert system can of course remember the resulting procedure and be able to repeat it without prompting if the user has a similar problem again.

Some expert systems are working now and others are being built. It appears that expert systems can (and probably will) be developed wherever the following criteria exist:

- The problem can be defined and solved in a logical, repeatable manner (if not, satisfactory computer programs are impossible).
- Problem definition and solution vary considerably over a range of situations, as in personal estate planning (if not, the computer's help is not needed—a book would do).
- The expertise involved has a specific scope, as in a professional practice area (if not, a specific knowledge base cannot be developed).
- It is worth spending $1 million to develop the system (this much will probably be needed before the system is in widespread use).

These constraints may seem discouraging, but we are looking ahead over decades. Each new system will add to the body of knowledge from those already existing, and each existing system will keep improving as more knowledge is added to it.

The mature, intelligent computer will be personalized and proactive. They will try to anticipate people's needs. Instead of complete publications, professionals and businessmen will be able to receive selected articles and items of information according to the interest profiles they create—which will be continually improved by the system without their knowledge as it observes their likes and dislikes. (This will not help everybody. If a user asks for different exception reports or varied formats every time, the system's efforts to anticipate may only impede communication.) News, entertainment, education and advice will all be subject to such customization and anticipation. Communications, too, will be subject to pre-stored interest profiles. The system might be informed, for example, that telephone calls from one's boss are to be put through at all times (and given an interrupt beep when one is already talking), but that all second-class mail is to be saved until after working hours.

Computers will provide guidance. Expert systems are likely to be developed in most activities now regarded as professional practice areas, and users of them will receive whatever degree of guidance the particular system is able to provide. Some functions, such as travel planning and directions, household repair, and shopping information can presumably be provided quite completely. Others, such as tax advice, estate planning, medical diagnosis and business modelling will presumably have limited service levels above which human experts take over.

Finally, the intelligent computer will be able to combine guidance and advice. If it has consulted with a user on his investment portfolio it will later be able to send him such messages as "A new bond has been issued which appears to meet your objectives better than the one you bought last month." The computer will also be able to intermix business, personal, educational and entertainment services. The dividing lines between business and personal activities may become more fluid.

Deciding when to turn from an expert system to a human expert will be a sensitive matter. How will one be able to tell when an expert system is giving subtly wrong advice that appears sound, but which will lead to long-term problems—or even legal liability? And how does one fix it? Tomorrow's intelligent computers will be asked to make inferences based on whatever knowledge they happen to have; not necessarily complete, correct, or up-to-date. They will, therefore, make errors. The mature intelligent computer will not be trustworthy.

It appears that information systems of the 21st century will always be what a child calls "dumb." They will only be able to do what they are told, and if the command is impossible or irrelevant the results will be too. The systems will try to identify and question destructive or unreasonable commands. They will even try to infer from the recent context that the user has made an error. But they will never understand what the user is trying to do; what he is "getting at." Since machines have no consciousness, volition or purpose they cannot generate a purpose that corresponds to that of the user. Intelligence has never existed without consciousness and self-awareness. Machines lack these, so artificial intelligence must always be fundamentally different from human intelligence—forever dumb.

Frederic Withington
Concord, Massachusetts