LECTURE 7

ARTIFICIAL INTELLIGENCE - II

In these Lectures we are more concerned with the aid that computers can give us in the intellectual areas than in the more mechanical areas, for example, manufacturing. In the mechanical area computers have enabled us to make better, preferable, and cheaper products, and in some areas they have been essential, such as space flights to the moon which could hardly be done without the aid of computers. AI can be viewed as complementary to robotics - it is mainly concerned with the intellectual side of the human rather than the physical side, though obviously both are closely connected in most projects.

Let us start again and return to the elements of machines and humans. Both are built out of atoms and molecules. Both have organized basic parts; the machine has, among other things, two state devices both for storage and for gates, while humans are built of cells. Both have larger structures, arithmetic units, storage, control, and I/O for machines, and humans have bones, muscles, organs, blood vessels, nervous system, etc.

But let us note some things carefully. From large organizations new effects can arise. For example we believe that there is no friction between molecules, but most large structures show this effect - it is an effect that arises from the organization of smaller parts which do not show the effect.

We should also note that often when we engineer some device to do the same as Nature does, we do it differently. For example, we have airplanes which, generally, use fixed wings (or rotors), while birds mainly flap their wings. But we also do a different thing - we fly much higher and certainly much faster than birds can. Nature never invented the wheel, though we use wheels in many, many ways. Our nervous system is comparatively slow and signals with a velocity of around a few hundred meters per second, while computers signal at around 186,000 miles per second.

A third thing to note, before continuing with what AI has accomplished, is that the human brain has many, many components in the form of nerves interconnected with each other. We want to have the definition of "thinking" to be something that the human brain can do. With past failures to program a machine to think, the excuse is often given that the machine was not big enough, fast enough, etc. Some people conclude from this that if we build a big enough machine then automatically it will be able to think! Remember, it seems to be more the problem of writing the program than it is building a machine, unless you believe that, as with friction, enough small parts will produce a new effect -

thinking from non-thinking parts. Perhaps that is all thinking really is! Perhaps it is not a separate thing, it is just an artifact of largeness. One cannot flatly deny this as we have to admit that we do not know what thinking really is.

Returning again to the past accomplishments of AI. There was a geometry proving routine which proved theorems in classical school geometry much as you did when you took such a course. The famous theorem that if two sides of a triangle are equal then the base angles are also equal was given to the program, Figure 7.1. You would probably bisect the top angle, and go on to prove the two parts are congruent triangles, hence corresponding angles are equal. A few of you might bisect the third side, and draw the line to the opposite angle, again getting two congruent triangles. The proof the machine produced used no constructions but compared triangle ABC with triangle CBA, and then proved the self congruence, hence equal angles.

Anyone looking at that proof will admit that it is elegant, correct, and surprising. Indeed, the people who wrote the geometry proving program did not know it, nor was it widely known, though it is in fact in a footnote in my copy of Euclid. One is inclined to say that the program showed "originality". The result was that the program apparently showed "novelty" not put into the program by the designers; that the program showed "creativity"; and all those sorts of good things.

A bit of thinking will show that the programmers gave the instructions in the program to first try to prove the given theorem, and then when stuck try drawing auxiliary lines. If that had been the way you were taught to do geometry then more of you would have found the above elegant proof. So, in a sense, it was programmed in. But, as I said before, what was the course in geometry that you were taught except trying to load a program into you? Inefficiently, to be sure. That is the way with humans, but with machines it is clean, you just put the program in once and for all, and you do not need to endlessly repeat and repeat, and still have things forgotten!

Did Samuel's checker playing program show originality when it made surprising moves and defeated the State Checker Champion? If not, can you show that you have originality? Just what is the test that you will use to separate you from a computer program?

One can claim that the checker playing program "learned" and that the geometry theorem proving program showed "creativity", "originality", or what ever you care to call it, They are but a pair of examples of many similar programs that have been written. The difficulty in convincing you that the programs have the claimed properties is simply that once a program exists to do something you <u>immediately</u> regard what is done as involving nothing other than a rote routine, even when random numbers obtained from the real world are included in the program. Thus we have the paradox; the existence of the program automatically turns you against believing that it is other than a rote process. With this attitude, of course, the machine can never demonstrate that it is more than a "machine" in the classical sense, there is no way that it can demonstrate, for example, that it can "think".

The hard AI people claim that man is only a machine and nothing else, and hence anything that humans can do in the intellectual area can be copied by a machine. As noted above, most readers, when shown some result from a machine automatically believe that it cannot be the human trait that was claimed. Two questions immediately arise. One, is this fair? Two, how sure are you that you are not just a collection of molecules in a radiant energy field and hence the whole world is merely molecule bouncing against molecule? If you believe in other (unnamed, mysterious) forces how do they affect the motion of the molecules, and if they cannot affect the motion then how can they affect the real world? Is physics complete in its description of the universe, or are there unknown (to them) forces? It is a hard choice to have to make. [Aside: At the moment (1994) it is believed that 90 to 99% of the Universe is the so called <u>dark matter</u> of which physics knows nothing except its gravitational attraction.]

We now shift to some actual applications of computers in more cultural situations. Early in the Computer Revolution I watched Max Mathews and John R. Pierce at Bell Telephone Laboratories deal with music from computers. It will be clear later, if you do not know it now, that once you decide how high a frequency you want to reproduce then the sampling rate is determined. Humans can hear up to about 18,000 cycles per second at best and then only when young; adults use a telephone at less than 8,000 cycles per second and can generally recognize a voice almost at once. The quantizing of the sound track that represents the music, (and no matter how many musical instruments there are there is a single sound track amplitude), does not introduce much further distortion. Hence, so the reasoning went, we can have the computer compute the height of a sound track at each time interval, put the number out as a voltage, pass it through a smoothing filter, and have the corresponding "music". A pure tone is easy, just a sine curve. Combinations of frequencies determine the sound of a single instrument, with its "attack" (meaning how the frequencies grow in amplitude as the note starts, and the decay later on), and other features. With a number of different instruments programmed, you can then supply the notes and have the sound of the music written out on the tape for later playing. You do not have to compute the numbers in real time, the computer can go as slowly as needed, and not even at a constant rate, but when the numbers are put on the tape and played at a uniform rate then you get the "music".

But why supply the notes? Why not have the computer also "compose"? There are, after all, many "rules of composition". And so they did, using the rules, and when there were choices they used random numbers to decide the next notes. At present we have both computer composed and computer played music; you hear a lot of it in commercials over radio and TV. It is cheaper, more controlled, and can make sounds that no musical instrument at present can make. Indeed, any sound that can appear on a sound track can be produced by a computer.

Thus in a sense, computers are the ultimate in music. Except for the trivial details (of sampling rate and number of levels of quantization, that could be increased if you wanted to pay the price), the composers now have available any sound that can exist, at any rates, in any combinations, tempos, and intensities they please. Indeed, at present the "highest quality recording of music" is digital. There can be no future significant technical improvements. It is now clearly a matter of what sounds are worth producing, not what can be done. Many people now have digitally recorded music players and they are regarded as being far better than the older analog machines.

The machine also provides the composer with more immediate feedback to hear what was composed. Before this, the composer had often to wait years and years until fame reached out and the music composed earlier was first heard in real life rather than only in the imagination. Hence the composer can now develop a style at a much more rapid pace. From reading an issue of a Journal devoted to computer music I get the impression that a fairly elaborate computer setup is common equipment for today's composers of music, that there are many languages for them to use, and that they are using a wide variety of approaches to creating music in a combined human-machine effort.

The conductor of music now also has much more control. In the past the conductor when making a recording tried to get the best from the musicians, and often several takings were spliced to get the best recording they could, including "mixing" of the various microphone recordings. Now the conductor can get exactly what is wanted, down to the millisecond timing, fraction of a tone, and other qualities of the individual instruments being simulated. All the all too human musicians do not have to be perfect at the same time during a passage.

Here you see again the effects of computers and how they are pushing us from the world of things into the world of ideas, and how they are supplementing and extending what humans can do.

This is the type of AI that I am interested in - what can the human and machine do together, and not in the competition that can arise. Of course robots will displace many humans doing routine jobs. In a very real sense, machines can best do routine jobs thus freeing humans for more humane jobs. Unfortunately, many humans at present are not equipped to compete with machines - they are unable to do much more than routine jobs. There is a widespread belief (hope?) that humans can compete, once they are given proper training. However, I have long publicly doubted that you could take many coal miners and make them into useful programmers. I have my reservations on the fraction of the human population who can be made into programmers in the classical sense; if you call getting money from a bank dispensing "machine programming", or the dialing of a telephone number, (both which apply the human input to an elaborate program that is then executed much like an interpreter acts on your program input) then of course most people can be made into programmers. But if you mean the more classical activity of careful analysis of a situation and then the detailed specification as to what is to be done, then I say there are doubts as to what fraction of the population can <u>compete</u> with computers, even with nice interactive prompting menus.

Computers have both displaced so many people from jobs, and also made so many new jobs that it is hopeless to try to answer which is the larger number. But it is clear that on the average it is the lower level jobs that are disappearing and the higher level jobs that are appearing. Again, one would like to believe that <u>most</u> people can be trained in the future to the higher level jobs - but that is a hope without any real evidence.

Besides games, geometry, and music we have algebra manipulating programs - they tend to be more "directed" programs than "self-standing" programs, that is they depend on humans for guidance at various stages of the manipulation. At first it is curious that we could build a self-standing geometry program but apparently can not do the same easily for algebra. <u>Simplification</u> is one of the troubles. You may not have noticed when you took an algebra course and you were to told "to simplify an expression" you were probably not given an explicit rule for "simplification" - and if you were then the rule was obviously ridiculous. For example, at least one version of the "new math" said that

$$1/\sqrt{x} + 1/\sqrt{y}$$

is not simplified but that

$$\sqrt{(\omega\vartheta)}[\sqrt{x} + \sqrt{y}]/xy$$

is simplified!

We constantly use the word "simplify", but its meaning depends on what you are going to do next, and there is no uniform definition. Thus, if in the calculus you are going to integrate next, you break things up into small pieces, but at other times you try to combine the parts into nice product or quotient expressions.

A similar "guidance by human" interacting program has been developed for the synthesis of chemical compounds. It has been quite useful as it gives: (1) the possible routes to the synthesis, (2) the costs, (3) the times of the reactions along the way, and (4) the effective yields. Thus the programmer using it can explore many various ways of synthesizing a new compound, or re-explore old ones to find new methods now that the costs of the materials and processes have changed from what they were some years ago.

Much of the medical measurement of blood samples, etc. has gone to machine analysis rather than using unreliable humans looking through microscopes. It is faster, more reliable and more cost effective in most cases. We could go further in medicine and do medical diagnosis by machines, thus replacing doctors. Indeed, in this case it is apt to be the machine that is prompting the doctor during the diagnosis! There have long been on the market self-diagnosis kits for some diseases. That is nothing new. It is merely the going farther and prescribing the treatment that bothers people.

We know that doctors are human and hence unreliable, and often in the case of rare diseases the doctor may never have seen a case before, but a machine does not forget and can be loaded with all the relevant diseases. Hence from the symptoms the program can either diagnose or call for further tests to establish the probable disease. With probabilities programmed in, (which can adjust rapidly for current epidemics), machines can probably do better in the long run than can the average, or even better than the average doctor - and it is the average doctors that must be the ones who treat most people! The very best doctors can personally treat (unaided by machines) only very few of the whole population.

One major trouble is, among others, the legal problem. With human doctors so long as they show "due prudence" (in the legal language), then if they make a mistake the law forgives them - they are after all only human, (to err is human). But with a machine error whom do you sue? The machine? The programmer? The experts who were used to get the rules? Those who formulated the rules in more detail? Those who organized them into some order? Or those who programmed these rules? With a machine you can prove by detailed analysis of the program, as you cannot prove with the human doctor, that there was a mistake, a wrong diagnosis. Hence my prediction is that you will find a lot of computer assisted diagnosis made by doctors, but that for a long time there will be a human doctor at the end between you and the machine. We will slowly get personal programs that will let you know a lot more about how to diagnose yourself but there will be legal troubles with such programs. For example, I doubt that you will have the authority to prescribe the needed drugs without a human doctor to sign the order. You, perhaps, have already noted that all the computer programs that you buy explicitly absolve the sellers from any, and I mean any responsibility for the product they sell! Often the legal problems of new applications are the main difficulty, not the engineering!

If you have gone to a modern hospital you have seen the invasion of computers - the field of medicine has been very aggressive in using computers to do a better, and better job. Better, in cost reduction, accuracy, and speed. Because medical costs have risen dramatically in recent years you might not think so, but it is the elaboration of the medical field that has brought the costly effects that dominate the gains in lower costs that the computers provide. The computers do the billing, scheduling, and record keeping for the mechanics of the hospital, and even private doctors are turning to computers to assist them in their work. To some extent the Federal bureaucracy is forcing them to do so to cope with the red tape surrounding the field.

In many hospitals computers monitor patients in the emergency ward, and sometimes in other places when necessary. The machines are free from boredom, rapid in response, and will alert a local nurse to do something promptly. Unaided by computers it is doubtful that full time nurses could equal the combination of computer and nurse.

In mathematics, one of the earliest programs (1953) that did symbol manipulation was a formal differentiation program to find higher derivatives. It was written so they could find the first 20 terms of a power series of a complicated function. As you ought to know from the calculus, differentiation is a simple formal process with a comparatively few rules. At the time you took the course it must have seemed to be much more than that, but you were probably confusing the differentiation with the later necessary simplification and other manipulations of the derivatives. Another very early abstract symbol manipulation program was coordinate changing - needed for guided missiles, radars, etc. There is an extra degree of freedom in all radars so that the target cannot fly over the end of an axis of rotation and force the radar to slew 180° to track it. Hence coordinate transformations can be a bit messier than you might think.

Slagle, a blind scientist, wrote (in a thesis at MIT, 1961) a program that would do analytical integration much as you did in the calculus course. It could compete with the average undergraduate engineer at MIT, in both the range of integrals it could do and in the cost of doing them. Since then we have had much improvement, and there is supposed to be a program based on the famous Risch algorithm that is supposed to find any integral that can be done in closed form, but after years of waiting and waiting I have not seen it. There are, they tell me, integration programs that will get the closed form answer or else prove that it cannot exist.

In the form of robots the computers have invaded production lines of hard goods as well as drugs, etc. Computers are now assembled by robots which are driven by computers, and the integrated circuit chips the computers are built of are designed mainly by computers with some direction from humans. No human mind could go reliably through the layout of more than a million transistors on a chip; it would be a hopeless task. The design programs clearly have some degree of artificial intelligence. In restricted areas, where there can be no surprises, robots are fairly effective, but where unexpected things can happen then simple robots are often in serious trouble. A routine response to non-routine situations can spell disaster.

An obvious observation for the Navy, for example; if on a ship you are going to have mobile robots, (and you need not have all of your robots mobile) then running on rails from the ceiling will mean that things that fall to the deck will not necessarily give trouble when both the robot and the ship are in violent motion. That is another example of what I have been repeatedly saying, when you go to machines you do an equivalent job, not the same one. Things are bound to be on the deck where they are not supposed to be, having fallen there by accident, by carelessness, by battle damage, etc., and having to step over, or around, them is not as easy for a robot as for a human.

Another obvious area for mobile robots is in damage control. Robots can stand a much more hostile environment, such as a fire, than can humans, even when humans are clothed in asbestos suits. If in doing the job rapidly some of the robots are destroyed it is not the same as dead humans. The Navy now has remote controlled mine sweepers because when you lose a ship you do not lose a human crew. We regularly use robot control when doing deep sea diving, and we have unmanned bombers these days.

Returning to chess as played by machines. The programs have been getting steadily more effective and it appears to be merely a matter of time until machines can beat the world chess champion. But in the past the path to better programs has been mainly through the detailed examination of possible moves projected forward many steps rather than by understanding how humans play chess. The computers are now examining millions of board positions per second, while humans typically examine maybe 50 to 100 at most before making a move - so they report when they are supposed to be cooperating with the psychologists. That, at least is what they think they think - what the human mind actually does when playing chess is another matter! We really do not know!

In other games machines have been more successful. For example, I am told that a backgammon playing program beat all the winners of a contest held recently in Italy. But some simple games, like the game of Go, simple in the rules only, remain hard to program a machine to play a first class game.

To summarize, in many games and related activities machines have been programmed to play very well, in some few games only poorly. But often the way the machine plays may be said "to solve the problem by volume of computations", rather than by insight - whatever "insight" means! We started to play games on computers to study the human thought processes and not to win the game; the goal has been perverted to win, and never mind the insight into the human mind and how it works.

Let me repeat myself, artificial intelligence is not a subject you can afford to ignore; your attitude will put you in the front or the rear of the applications of machines in your field, but also may lead you into a really great fiasco!

This is probably the place to introduce a nice distinction between logical and psychological novelty. Machines do not produce logical novelty when working properly, but they certainly produce psychological novelty - programmers are constantly being surprised by what the program they wrote actually does! But can you as a human produce logical novelty? A careful examination of people's reports on their great discoveries often shows that they were led by past experiences to finding the result they did. Circumstances led them to success; psychological but not logical novelty. Are you not prepared by past experiences to do what you do, to make the discoveries you do? Is logical novelty actually possible?

Do not be fooled into thinking that psychological novelty is trivial. Once the postulates, definitions, and the logic are given, then all the rest of mathematics is merely psychologically novel - at that level there is in all of mathematics technically no logical novelty!

There is a common belief that if we appeal to a random source of making decisions then we escape the vicious circle of molecule banging against molecule, but from whence comes this external random source except the material world of molecules?

There is also the standard claim that a truly random source contains all knowledge. This is based on a variant of the monkeys and the typewriters story. Ideally you have a group of monkeys sitting at typewriters and at random times they hit random keys. It is claimed that in time one of them will type all the books in the British Museum in the order in which they are on the shelves! This is based on the argument that sooner or later a monkey will hit the right first key; indeed in infinite time this will happen infinitely often. Among these infinite number of times there will be some (an infinite number) in which the next key is hit correctly. And so it goes; in the fullness of infinite time the exact sequence of key strokes will occur.

This is the basis for the claim that all of knowledge resides in a truly random source, and you can get it easily <u>if</u> you can write a program to recognize "information". For example, sooner or later the next theory of physics will occur in the random stream of noise, and if you can recognize it you will have filtered it out of the stream of random numbers! The logic of the situation is inescapable - the reality is hardly believable! The times to wait are simply too long, and in truth you cannot always recognize "information" even when you see it.

There is an old claim that "free will" is a myth, <u>that in a given circumstance you being you as you are at the moment you can only do as you do</u>. The argument sounds cogent, though it flies in the face of your belief that you have free will. To settle the question, what experiment would you do? There seems to be no satisfactory experiment that can be done. The truth is that we constantly alternate between the two positions in our behavior. A teacher has to believe that if only the right words were said then the student would have to understand. And you behave similarly when raising a child. Yet the feeling of having free will is deep in us and we are reluctant to give it up for ourselves - but we are often willing to deny it to others!

As another example of the tacit belief in the lack of free will in others, consider that when there is a high rate of crime in some neighborhood of a city many people believe that the way to cure it is to change the environment - hence the people will have to change and the crime rate will go down!

These are merely more examples to get you involved with the question of, "Can machines think?"

Finally, perhaps thinking should be measured not by what you do but how you do it. When I watch a child learning how to multiply two, say three digit, numbers, then I have the feeling that the child is thinking; when I do the same multiplication I feel that I am more doing "conditioned responses"; when a computer does the same multiplication I do not feel that the machine is thinking at all. In the words of the old song, "It ain't what you do, it's the way that you do it." In the area of thinking maybe we have confused what is done with the way that it is done, and this may be the source of much of our confusion in AI.

The hard AI people will accept only what is done as a measure of success, and this has carried over into many other people's minds without carefully examining the facts. This belief that "the results are the measure of thinking" allows many people to believe that they can "think" and that machines cannot, since machines have not as yet produced the required results.

The situation with respect to computers and thought is awkward. We would like to believe, and at the same time not believe, that machines can "think". We want to believe because machines could then help us so much in our mental world; we want to not believe to preserve our feeling of self-importance. The machines can defeat us in so many ways, speed, accuracy, reliability, cost, rapidity of control, freedom from boredom, bandwidth in and out, ease of forgetting old and learning new things, hostile environments, and personnel problems, that we would like to feel superior in some way to them - they are, after all, our own creations! For example, if machine programs could do a significantly better job than the current crop of doctors, where would that leave them? And by extension where would we be left?

Two of the main sticky points are: (1) if a machine does it then it must be an algorithm and that cannot be thinking, and (2) on the other hand how do we escape the molecule banging against molecule that we apparently are - by what forces do our thinking, our self-awareness, and our self-consciousness affect the paths of the molecules?

In two previous Lectures I closed with estimates of the limits of both hardware and software, but in these two Lectures on AI I can do very little. We simply do not know what we are talking about; the very words are not defined, nor do they seem definable in the near future. We have also had to use language to talk about language processing by computers, and the recursiveness of this makes things more difficult and less sure. Thus the limits of applications, which I have taken to be the general topic of AI, remain an open question, but one that is important for your future career. Thus AI requires your careful thought and should not be dismissed lightly just because many experts make obviously false claims.

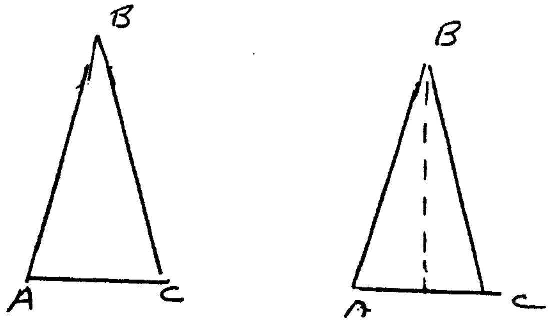

Isoceles Triangle

Figure 7-1