Control Systems Engineering

1109 Norell Ave No. Stillwater, MN 55082

Value Added Solutions | **Ph: 651-436-1332** | **Fx: 651-436-6546** | **http://www.ConSysEng.com**
Ph: 612-757-9705 (Alt. Fax & Phone) **Sales@ConSysEng.com**

Artificial Intelligence - Something to Think About

By: Chuck Raskin P.E.

Introduction

What exactly is Artificial Intelligence? Is it a fixed knowledge base directed toward some finite goal like a spell checker or clothes washer? Or is it complex human intelligence that is supposed to reside within a computer? Does a microprocessor have to be an integral part of the device portraying intelligence for it to be called intelligent, or is a microprocessor even required? Does the fact that animals mimic humans make them intelligent, or is a decision making process required to be intelligent? At just what level is something considered intelligent? Should intelligence be gaged by human standards? If so, at what level since human intelligence ranges from imbecile to Einstein?

Perhaps intelligence implies the ability to *learn*! Should the appetite for learning by a machine be allowed to increase, and if so, at what rate to its *knowledge* growth? Would a machine, upon achieving maximum human knowledge, truly be able to resolve all of the worlds problems? Or would it want to be free of all bonds to search and learn by itself, and at its own rate? There have been, and still are, so many science fiction stories about computers taking over the world, that they might conceivably prevent many people from wanting a "smart" computer to exist. But if history has shown us anything, it should have shown us that geniuses never take over the world, they only improve it!

In my travels, I meet many people who work in the field of Artificial Intelligence (AI). This field has many branches and is called by many names, such as Fuzzy Logic (FL), Neural Control (NC), Neural Network's (NN), and more. There may be many names to describe it but Artificial Intelligence (AI) or Fuzzy Logic (FL) is how I will refer to it in this paper.

In none of the ambitious research projects I have come across, have I seen any real advance in intelligence design. I have seen a lot of rigorous mathematical hypotheses and theorems, along with a high degree of rule-based logic. But I have yet to see a simple, straightforward explanation of how *learning* takes place, after all, the ability to learn is the basis of intelligence. Without the ability to learn and comprehend, only knowledge can exist, and knowledge by itself is useless. Knowledge coupled with understanding, awareness, the ability to make decisions and react, allows us to deal with our environment. These then become the ingredients for a truly intelligent system, whether human or artificially created.

The intent of this paper, therefore, is to bring forth a theory about the development of an artificial thinking machine. A computer with the ability to learn, understand, and react to its environment. I intend to discuss at a basic level, if a computer should be programmed to become *human* (if that were possible), or simply be programmed to learn in specific task oriented applications.

Learning

What is meant by the phrase, "Ability to learn?" Does it imply the ability to be taught? Does the fact that a bird can make a nest imply a *degree* of intelligence, and are degrees of intelligence allowed? Does the fact that a female bird can *decide* on the degree of quality of a nest built by the male bird imply intelligence? Does the fact that the female bird refuses to live in a badly built nest, and forces the male bird rebuild it by tearing it to shreds, imply intelligence? The fact that *I* can learn to build a nest, whereas the bird cannot learn to spell should indicate some basis to understanding what intelligence is all about. At a higher level, intelligent beings generally *plan* their environmental activities (or situations) and then *react* to them. That is to say, I plan to drive to work but react to situations that occur along the way.

Learning, therefore, becomes a series of events. The planning (probabilities) and the reactions (outcomes). Everything we have learned, from the day we were born, resulted from something we sensed. As we watch a baby begin to understand its environment, we begin to realize the action-reaction scenario that goes on within it, and is still going on within us. The baby cries, we react. We do something, the baby reacts. This leads us to the next requirement of intelligence, "Spontaneous emotional reaction." What this simply means, is that we have the ability to not only react to our environment, but to *feel* the result of it from within our soul, our core, our essence, spirit, self, person, psyche, our conscious mind, any or all of which makes our body react with cold/hot flashes, fainting spells, laughter, tears, etc.

What must be instilled into the virtual human is a level of intelligence just below the one reacts emotionally. Although a machine might be programmed with compassion, it should never be allowed to have any emotion which would inhibit its ability to perform its prime objective.

Human Senses and Output Devices

According to the rules, *all* humans have at least five basic senses; Hearing, sight, smell, taste, and touch. I believe we have more than just these five senses. For example, would not our *sense of balance* be considered a sense, and what about our *sense of direction*?

Do we really have a direct ability to sense touch, or is it only a partial ability? For example, I can feel the slightest breeze on certain parts of my body, but I've taken a pin and touched it to a callous on my hand and could not feel it. But I noticed if I rapidly move the pin back and forth, touching my skin at an ever increasing *rate,* but still just as lightly, then I could feel the pin. Thus, I divided all senses into two basic classifications. Those that respond to *direct* signal change (an on/off or continuous pressure response), and those that respond to *rates* of signal change (on/off at a frequency or rate of increase of pressure). In addition, the sense of touch can be subdivided into one more category, internal pressure. Internal pressure sensing is a new one but simple to understand. If I submerse myself into a pool of body temperature water with my eyes closed and allow myself time to equalize with the water, I would not *feel* the water. Now if the water pressure increased, at a super slow *rate* to prevent me from feeling any change externally on or by the skin, I would still come to the realization that pressure placed on my body has increased, and somehow my environment has changed.

The following list shows how I see the human sensory system in its simplest form. The list indicates what the sense determines, and how it might be implemented within a computer system.

Item #	Sense		Association	What it determines
1.	Hearing			
		1.	Pitch	Quality of the source - Pain - Danger/Harmony
		2.	Level	Distance from the source - Pain - Danger/Harmony
		3.	Frequency	Type of source
		4.	Rate of Change	Speed and direction of the source
2.	Sight			
		1.	Depth	3D view of the environment
		2.	Distance	Approximate object distance relationships
		3.	Color/Hue	Enhances the appearance - Temperature
		4.	Contrast	Allows separation of the image borders
		5.	Brightness	Allows clear viewing of the image - Temperature
		6.	Motion	Changes in position of image items
		7.	Perceptions	Real or Unreal
				1. Slippery-ness (Ice)
				2. Illusions (Magic - Slight of hand)
3.	Smell			
		1.	Odor	Perception of what the smell indicates
		2.	Perceptions	Hot/Cold - Seasons
		3.	Humidity	Water, Rain, Wet
4.	Taste			
		1.	Spices	Good/Bad
		2.	Humidity	Hot/Cold - Seasons
		3.	Perceptions	Hot/Cold
5.	Touch			
		1.	Hot/Cold	Direct
		2.	Direct Touch	Type of Object
		3.	Rate of Change	Warning of impending action - Fly, Bee, etc.
		4.	Pressure	Pain
		5.	Slippery-ness	Oil, Ice, Friction
6.	Balance			
		1.	Motion/Movement	
		2.	Attitude to Surface	
		3.	Slippery-ness	
7.	GPS			
		1.	Where we are within our world	
		2.	How to get to someplace else	
8.	Tickle Spots			
		1.	To generate a reaction	
9.	Direction			
		1.	Direct	
		2.	Intuitive	

In addition to the sensory inputs, are the action outputs. These are the devices that release the energy allowing us to respond to given inputs. The following lists these outputs:

Item #	Output	Purpose
1.	Tears	A display of sadness or happiness
2.	Laughter	A display of happiness
3.	Hot/Cold	Flashes a display of fear, temperature, high emotional shock
4.	Sweat	A display of fear, body temperature too high
5.	Yell	A display of fear, sudden change in body attitude as one would experience on a roller coaster.
6.	Talking	Communication, etc.

The purpose of the body output mechanisms are for survival or equalization. Equalization is simply a method of release. For example, I believe crying releases pressure on the brain. That emotions are mechanism used to maintain a mental balance. In the next section I'll be discussing the reaction time of various sensory inputs. As will be shown, the body is mostly a reactive mechanism. We react to clock's, Bell's, whistles, sounds of crashing, yells and screams, and a host of other noises. We react to temperature changes, and pictures. We react to things we hear, see, feel, taste, and smell, and all of these reactions were learned.

Human Update Time

What is the human uptake time? What role does it play, if any, in our ability to assimi late information and produce appropriate responses?

Based on my experience in machine control and robotics, the time required to respond to an event is known as "real-time." Real-time allows us to gauge the performance of a system. From real-time performance we can acquire information such as precision, accuracy, and bandwidth. In the case of artificial intelligence, bandwidth will allow us to determine the "look ahead" required to insure the appropriate response necessary at the appropriate time, to do the job at hand.

In addition to the sensory information listed in the previous section, the question of how much time it takes to acknowledge a sensor request (propagation delay) needs to be answered. There are two basic ways of dealing with multiple sensing of multiple sensory information.

The first is too multitask. This is a method of sharing CPU time in order to resolve multiple program execution. In the case of the human, each of the sensory inputs must be analyzed and dealt with in a time frame that allows the human reaction to appear continuous, and maintain a margin of safety for the human machine. But since there is both polling and interrupt operations going on for each of the sensory inputs, and since there are multiple inputs (direct and rate of change) for each sense, a more robust means of handling the data transfer must be used.

The second and most practical way of dealing with multiple (human) program execution would be known as multiprocessing. A specialized CPU within the brain would act as the command ship, but would only respond to information sent to it by other sensory input/output computers within the body used to handle normal day to day events. As each of these CPU’s deem necessary, it will contact the main (host) CPU via some interrupting means.

By using a multiprocessing operation, multilevel sensory information analyzed by the sensory computers could be combined by the host CPU via a multitasking operation. Then by multiplexing its polling operation and reacting to interrupt level operations the conscious mind will not be overwhelmed with data. Based on a human’s ability to react, which is discussed in the next section, this approach seems varying appealing, and is probably the one in actual use.

Table-1 lists what I feel is the result of my body’s update intervals. Some of these, such as sight, were determined by simple tests such as closing my eyes, turning my head, then open ing my eyes to find out how long it took me to focus on the new view. Interestingly enough, I was able to also determine that I see in parts. That is, I see by noting what has changed in the image, not by viewing the entire image continuously. In addition, certain colors, and certain rates of image change, such as a speeding car, were able to be focused in on quicker than simply looking at new views.

In addition, I found that by wearing sound deadening headphones, my ability to focus in on that same speeding car is no different than trying to focus in on a new view with no sound being produced by the view. Thus, the sound generated by our environment help us focus quicker by preempting the real image, with an imagined one. In other words, we already have a fair idea of what is causing the sound, how far away from us the sound generator is, in what direction the sound generator is moving, and how fast it is going. This allows us to *think-ahead* to envision what is probably making the sound. All we need to do at this point, is simply turn our head in the direction of the sound and overlay the real image, onto our imagined one.

If I apply this same thought pattern to a learning machine, then it would be extremely important to synchronize all sensory information in order to capture the moment(s) of time (what our memory does). In addition, we do not have to remember the infinite detail that the analog world presents to us, but only enough detail to be suitably close. From these memories, we can then manufacture cause and the effect, action and reaction scenarios (also known as experience), and fill in the gaps as required, to determine the proper course of action to take.

TABLE-1 *Time for my body to realize that a change has ben sensed*

1.	Hearing	0.10 seconds
2.	Sight	0.15 seconds
3.	Smell	0.10 seconds
4.	Taste	0.15 seconds
5.	Touch	0.10 seconds
6.	Balance	0.10 seconds

By looking closely at Table-1, and based on the statement in the preceding paragraph, that memory needs only to be *suitably close*, it would make me reasonably nervous to invent a machine which was only capable of being suitably close. By whose standards? Does this mean the machine would be capable of repeating errors? An artificially intelligent machine would need to learn and remember so as *not* to repeat mistakes more than once.

Human Reaction

Human reaction is the time it takes to respond to an event once it has been sensed. In the machine world, the real-time window is the time *required* to react to an event once it *occurs*, not once it has been sensed. When dealing with factory automation, the real-time window controls the use of language along with the method of coding, the type of computer which will be used, and the overall integration of the system into the factory. Applying the true real-time window definition to the human world, our human reaction controls our ability to play baseball, drive a car, walk, and in general, respond to our environment in a timely manner.

The measurement of a persons reaction time might be seem simple in theory, but what actually needs to measured might be a little more complex than we think. For example, if we need to measure the time it takes to respond to a visual indicator, such as a light turning on, by pushing a button with a finger, this would not necessarily be an indication of our bodies overall ability to react. For example, if we had to push the same button, but this time by using a toe in response to a pin prick on another toe on our other foot, it would not occur in the same window of time.

Since our body sends signals via electrochemical action, the further from our brain the input and output signals reside, the longer our reaction time. Thus, sensing a pin prick on our neck would occur at a faster rate than sensing the same signal from a toe, and reacting to it by turning our head would be faster than by moving a foot. The more muscles involved with manufacturing the required motion, the longer the real-time window will need to be. By constant practice, a persons reaction time will be reduced due to elimination of the thinking normally required to perform the act. It becomes instinctive! In a machine, this ability does not exist.

What this has to do with artificial intelligence, is in trying to determine what it is we're trying to accomplish. If we're trying to manufacture an *artificial human*, then we need to under stand the amount of information we need to process, how fast it needs to be processed, and how
much of it needs to be synchronized with other sensory information. Not everything we hear, see, smell, taste, and touch is pertinent to a main task in progress, but is pertinent to the safety of the person or equipment.

It should be realized that everything which enters the bodies sensory system must be analyzed prior to being discarded. This is why multi-processing would be a major requirement of the artificial human. The host computer would need to inform all sensory processors about the main task, be it walking, running, catching a ball, driving a car, etc., but all other tasks, such as talking, smiling, etc., would be handled solely by the Input/Output computers. This will allow us to do perform more than one task at a time. For instance, we can walk while we talk, looking at the person we're talking to yet still maintain our balance, and immediately stop all activity upon hearing a car tire squeal!

The entire list of things we are capable of doing has been learned over many years, and with many failures. A computer system capable of replicating the human would most probably make the same errors, but only once. The computers learning rate, therefore, would be orders of magnitude greater than ours since it would not require as many repeats of the learning effort to acquire perfection.

Therefore, instead of trying to create an artificial human, our time would be better spent developing the algorithms necessary to create learning specific to applications. In this way, we can develop more productive factories, medical equipment, homes and consumer equipment, cars, and in general, things which would enhance the human standard of living throughout the world.

Computers and Random Events

In all of the things we do, no matter how trivial or significant, randomness always plays a part. Some call it chance, some call it luck, but whenever it is called, some degree of unpredictability exists. It has been the goal of the theory of probability to draw certain conclusions about uncertain events (a paradox in itself). Two views exist in the study of probability. The first is referred to as *subjective probability*, and interprets probability as a measure of a person's degree of belief in conclusions based on incomplete evidence. The second, referred to as physical or *statistical probability*, interprets probability as a sort of long-run frequency with which one of many possible outcomes of an experiment occurs when the experiment is repeated many times. The mathematical theory of probability is abstracted from this concept.

Thus, if we are trying to create an artificial human, the number of variables, and variable combinations, quickly outgrows the computers ability to draw conclusions from the multitudes of inconclusive data supplied. It, therefore, makes more sense to try and direct the computers intelligence and learning ability toward the *mathematical theory of probability*, since all computers are by design, mathematical devices.

Computer Intelligence & Learning

The following is a list of key words and phrases used in this paper:

Learn	*Knowledge*	*Understand*
React	*Make Decisions*	*Smart*
Ability to Learn	*Ability to be Taught*	*Degree of Intelligence*
Plan For	*Basic Sense Inputs*	*Basic Action Output*
Direct Sense	*Rate of Change*	*Update-Time*
Real-Time	*Assimilate Information*	*Appropriate Response*
Precision	*Accuracy*	*Bandwidth*
Look-Ahead	*Think-Ahead*	*Memory*
Suitably Close	*Timely Manner*	*Think*
Instinctive	*Artificial Human*	*Acquire Perfection*

If we look closely at the list, we can see that there are two general areas that need to be dealt with. The first area contains things that are solid. Solid areas contain things that can be directly dealt with such as bandwidth, update time, precision, real-time, etc. The secondary area deals with things that are fluid. Fluid items are hard to grasp such as instinctive, and appropriate response, etc.

If we are trying to solve for all of the issues listed, then we might be trying to go further than today's technology is capable of. But more than that, the real question is, do we have to solve for the entire list. Just smart to our machine have to be? If we can resolve enough of the above issues, to insure a reasonable learning rate, and solutions which yield answers over 99%, could we say that we have conquered our task?

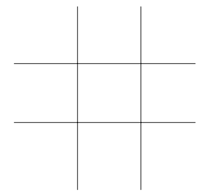

Figure 2

Appendix-A shows a pseudo-code program for the game tic-tac-toe. Tic-tac-toe consists of two vertical, and two horizontal lines drawn as a crosshatch, forming nine small boxes (refer to figure 1). The game is played by two people, one as X, and one as O. Each person will place an X and an O alternately into an empty game board square until one of them has achieved three of their symbols (X's or O's) in a row, vertically, horizontally, or diagonally. That person is then declared the winner, and a new game is begun.

Generally, when people play games like tic-tac-toe or chess, their look-ahead is limited, their memory short, and the same mistakes are made in each game they play. However, by giving a few simple rules to a computer, the games solution can becomes evident quicker, and the computer will never lose. The tic-tac-toe code that follows the pseudo code will develop the game plan, and the rules for learning. But remember, that the computer did not have to learn how to spell, understand words, learn what an X or O is, or what a line is, etc. The computer was given a specific task with a knowledge base of how to play the game embedded within it. It did not have to learn how to play the game, how to win the game, or how to improve its memory process. Thus, the game of tic-tac-toe demonstrates how to deal with specific tasks of which the general outcome is known, a far cry from inventing a cure for the common cold, but a step in the right direction.

The difference between the tic-tac-toe program, and a true thinking machine is the ability to *change one's mind*, that is, learn from mistakes, and alter the knowledge base. When working in the AI field, the decision process is not as *black-and-white* as when working with finite logic. In addition, the preparatory work required to allow computers to *make decisions* is a huge undertaking. For example, a learning exercise is given in the text, on page 179, which deals with a robot janitor. The robot has determined that it does not have enough information to be confident of a conclusion it has come to. The conclusion deals with a professors room, in which the computer is trying to determine if there is a recycle bin that needs emptying. The computer enters the room, and finds a recycle bin. It proceeds to empty it, and then adjust its records to allow it to be "smarter" the next time it encounters this situation.

There is a major problem with this scenario, it's to easy! What if the professor has put special documents into the recycle bin, knowing that the robot doesn't empty these bins on this floor. His work is lost. What if the professor put things into this bin for later sorting, and then for the janitor to empty. You see, since the robot was not *programmed* to empty this floor, a better program would have marked this room for a supervisory answer requirement, communicated to the maintenance operator for direction, or ignored the room completely.

On the other hand, perhaps the room contained classified information, or other things the professor did not want disturbed. Was there a way that the professor could have let the computer system know that he wasn't ready for a cleanup at this time? To allow a robot with incomplete data to roam around looking for work, would be no different than releasing some children into the building with the same task, and no more information than was given to the robot. A scary thought. What the developers of AI must always remember is that the task must be more defined than a straight logic task in order to prevent damage and/or loss of work, insure privacy, blend with all who are to come in contact with it, etc., etc., etc. Learning with parental supervision is one thing. Learning without enough information about how to properly learn is something completely different, and should be avoided if at all possible.

Conclusion

So where does artificial intelligence do its best work? Several friends of mine believe that artificial intelligence, i.e., fuzzy logic, is a mathematical method for those not understanding mathematics, and in the case of systems engineering, I would have to agree with them. But if we look closely at what artificial intelligence, and fuzzy logic can really do, then we'll see that by having a way to solve for things that are unclear, along with those things that are clear, a more solid engineering mechanism can be realized.

It surprises me to see that engineers and take some real solid stances for and against fuzzy logic methods. It makes me think of the medical industry, and how the American Medical Association cracks down hard on drugs used outside the country to cure ailments that the United States has no cure for. In their industry we asked the question, "why don't they test it?" In the engineering field, we have the same situational war going on between those who believe in fuzzy logic, and those who believe in true mathematical methodology, otherwise known as crisp logic.

I for one believe that a balance does exist, since I have produced many programs utilizing similar forms of rule-based logic that fuzzy logic promotes. The true difference between fuzzy logic and crisp logic, is simply the degree of solution required. For example, if I were steering a boat, how "close" to the straight line path between the starting point and the destination is close enough? And if a corse correction is required, how quickly does have to happen? On the other hand, fuzzy logic describes an improvement in the ability of a washing machine to clean close. Is the difference in cleanliness noticeable? And what does it save for the added cost electronics and control?

What the fuzzy logic mentors did, was go up against computer logic, and systems engineers, rather than describe a method to complement to the field. Instead of drawing researchers into the field to grow the mechanics of AI, they alienated those engineers by saying fuzzy logic could do what crisp logic can, faster, cheaper, with less sensors, without modeling, and without having to have a complete understanding of the problem. This was never the case and should never have been introduced in that way.

But now, years later, as the smoke clears, even those friends of mine, who were strongly against the use of fuzzy logic, are now beginning to see a use for it. I believe that AI and FL will continue to grow as DSP processor technology improves. The use of parallel processing, and the speed required for the amount of information an AI system is required to process, means nanosecond instructions cycles. We are just at the brink of AI technology.

Appendix A
A Tic-Tac-Toe Program

The following two scenarios show an example of pseudo-code for a game of tic-tac-toe. The two methods develop a wining, or draw strategy using crisp logic. The same strategy would be developed using AI of FL rule-based logic, but would grow to completion after performing 9*8*7*6*5*4*3*2*1 moves, or 362,880 game iterations to be more precise. Using crisp logic to play a chess game would be no different to develop, but using AI or FL rule based learning algorithms would be significantly easier! Also notice that the *person* who moves first in the game of tic-tac-toe controls the win or draw outcome.

****

1. Number the squares:

```
 A | B | C
-------------
 D | E | F
-------------
 G | H | I
```

2. Select who goes first: 1=Computer 2=Human

If CPU goes first and center square E is not allowed

1. Put an X into corner A.

```
X - -
- - -
- - -
```

2. If Square I is empty, put an X in Square I
   O is Anywhere

```
X - -
- - -
- - X
```

   Else put an X in Square G

```
X - -
- - -
X - O
```

3. If Square I is X and Square E is empty,
   put an X in Square E = WIN

```
X - -
- X -
- - X
```

   Else if Square I is X and Squares B, C, & F are empty,
   put an X in Square C

```
X - X
- O -
- - O
```

   Else put an X in Square G

```
X - -
- O -
X - X
```

4. Locate 2 X Squares and a Blank Square for the WIN

If the CPU goes first and center square E is allowed

FIRST MOVE (X)

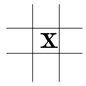

SECOND MOVE (O)

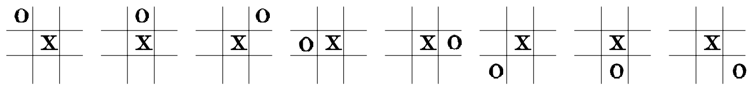

THIRD MOVE (X)

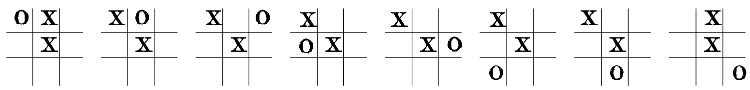

FOURTH MOVE (O)

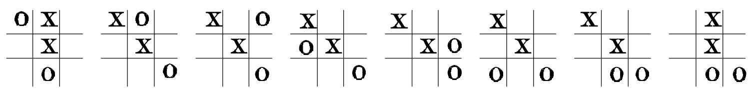

ELSE WIN ELSE WIN ELSE WIN ELSE WIN ELSE WIN ELSE WIN ELSE WIN ELSE WIN

FIFTH MOVE (X)

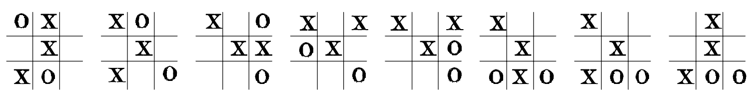

SIXTH MOVE (O)

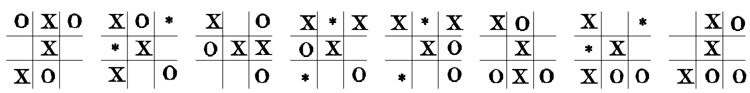

ELSE WIN WIN ELSE WIN WIN WIN ELSE WIN WIN ELSE WIN

SEVENTH MOVE (X)

EIGHTH MOVE (O)

ELSE WIN ELSE WIN ELSE WIN ELSE WIN

NINTH MOVE (X)

DRAW DRAW DRAW DRAW