Institute for Christian Teaching
Education Department of Seventh-day Adventists

A Christian Perspective on the Study of Computational Intelligence

by
James Wolfer
Department of Computer Science
Andrews University
Berrien Springs, MI U.S.A.
email: wolfer@andrews.edu

Prepared for the
International Faith and Learning Seminar
held at
Newbold College, Bracknell, Berks, England
June 1994

Abstract

The study of Machine Intelligence is currently in a state of transition, with new paradigms emerging which challenge both the Christian student and educator. Research topics which would have been relegated to the fringe of the scientific community twenty years ago are now being pursued in the mainstream. Aggressive research is now conducted in such areas as artificial neural networks, genetic programming, and artificial life. These computational models are based on abstractions from nature, including human neural anatomy, Darwinian evolution, and the basic assumption that life, in general, can be reduced to information. This includes the possibility that eventually machines will exceed the intellectual capacity of human beings, and that humans may be creating their own evolutionary replacements.

The ultimate fate of these research efforts is, of course, an open question. Potential limitations of Artificial Intelligence have been discussed during previous seminars in this series. Independent of the outcome, however, are serious moral and ethical concerns. If we believe that we are creating our successors, then, in the limit, we encounter the risk of considering human life an obstacle to that goal and therefore expendable.

Within this framework we, as Christian educators, must give our students a place to stand. This may be one of the meaningful contributions that we as Christians can provide to mankind. Since we define what it means to be human in terms of a transcendent relationship, independent of our relative intellectual capacity, we have a secure platform from which to study artificial intellect. And, since our worth as human beings is derived from God's demonstration of our worth to him, we have a stable foundation for evaluating ethical and moral concerns independent of human goals.

1 Introduction

Perhaps the best way to introduce the profound changes facing society at large, and computer science students specifically, is to examine a book review by Hans Moravec, of the Carnegie Mellon University Robotics Institute [1]. Reviewing the book Metaman [2], which examines the potential of society itself evolving into an independent organism, eventually transcending human form and intellect, Moravic remarks:

> ...He (Stock) suggests early on that human beings are likely to be a part of Metaman indefinitely, but later notes there are technologies that will probably totally reshape – or replace – humans. In my opinion, he greatly overstates the long-term importance of the human form. Metaman evolves so quickly that essentially everything we know, including ourselves, is in the process of becoming history...We can hope for a comfortable retirement, or help in restructuring ourselves into something more useful, but our current bodies and minds will be increasingly anachronistic.

After this summary, Moravic concludes,

> Metaman is a well-written book whose point of view, if more widespread, would reduce the number of frightened people angrily tilting at the windmills of a rapidly changing world.

Statements such as these, along with the pursuit of artificially cognizant machines, evolutionary programming, and neurally inspired computer architectures form the context for much of the study of computational intelligence within the last few years. In this paper I will explore these emerging paradigms, place them within the context of traditional Artificial Intelligence, and explore some of their implications from a Christian perspective.

2 Current Directions in Machine Intelligence

2.1 Classical Artificial Intelligence

Classical Artificial Intelligence (AI) has attempted to apply formal computational procedures to problems requiring intelligence when performed by human beings. By applying techniques such as formal logic, computability theory, and various ad-hoc methods (e.g. search) traditional AI has tried to reproduce the high-level functionality observed in intelligent creatures without necessarily modeling the computational mechanism they possess. While traditional AI has provided some success, such as expert systems, game playing strategies, and a formal mathematical foundation it has essentially failed on the day-to-day survival strategies, such as vision, that even the simplest creatures perform trivially. These insufficiencies have given rise to a new series of paradigms which attempt to model not only the behavior, but also the process which is assumed to have produced the behavior. Since the subject of the strengths and limits of traditional Artificial Intelligence has been addressed during previous seminars in this series [3] it will not be repeated here.

In the following subsections some of these new paradigms will be described, followed by a discussion of some potential social implications.

2.2 Artificial Neural Networks

Neural networks can be viewed as an instance of a set of paradigms which are sometimes characterized as emergent computation[4]. These techniques share the concept that solutions to complex problems can emerge from simpler adaptive computations. Efforts in this area (including neural networks) often attempt to mimic some elementary aspects of natural phenomenon, natural selection in the case of genetic algorithms and neural physiology in the case of neural networks. To a large extent, these are metaphorical relationships, since our current ability to model extensive networks is computationally

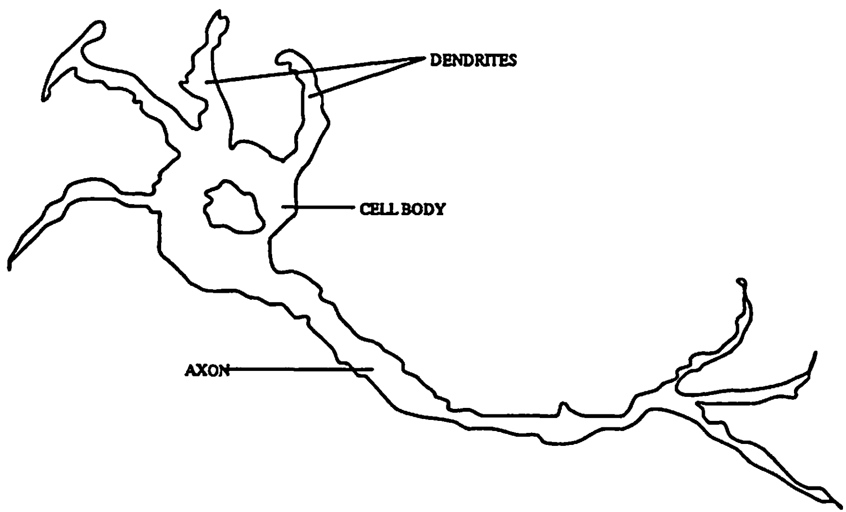

Figure 1: Neuron Schematic Diagram [5]

limited. An examination of the neural ability of the human brain serves to illustrate this point.

Figure 1 schematically illustrates a typical neuron [5]. The major structures to note are the cell body, the dendrites, and the axon. The axon serves the transmission function, while the dendrites act as receptors from other neurons. The intersection between an axon branch and the receiving neuron is called a synapse. Electrical signals are transmitted by the repeated chemical depolarization of the neural tissue. It is believed that information is transmitted in the form of pulse coded analog signals, which have either an excitatory or an inhibitory effect on the receiving neurons.

While possessing complex chemistry, these cells form a fairly primitive processing element. To compensate, natural systems possess extremely large numbers of neurons. Current estimates of the neural density in the human cerebral cortex range from 10^{10} to 10^{12} neurons, providing about 10^{15} synaptic connections working in parallel. This contrasts with the small number of neurons in artificial neural simulation.

Artificial neurons mimic their natural equivalents by substituting a matrix representing the strength of the interconnections for the pulse frequency encoding. Information is encoded and distributed throughout the network in terms of these weights. Training the network is the process of obtaining a viable set of weights for the problem at hand.

There are many varieties of artificial neural networks. These are characterized by their connection architecture, training strategies, and feedback mechanisms. Kosko [6] provides the general taxonomy which is abridged in Figure 2.

In feedforward networks data is presented to the inputs and the result of the com-

Feedforward Supervised	Feedback Supervised
Backpropagation	Recurrent Backpropagation
Feedforward Unsupervised	**Feedback Unsupervised**
Self-Organizing Maps Competitive Learning	Art-1 Art-2

Figure 2: Neural Network Taxonomy [6]

putation is produced at the outputs. The inputs have no information about current or previous outputs. Feedback networks, on the other hand, make current (or time-delayed) output available to the inputs for the next computational cycle.

Supervised training presents a set of inputs and the corresponding target outputs to the network. The training process then attempts to generate interconnection weights which cause the network to produce the target output corresponding to each set of inputs. Unsupervised training, on the other hand, presents only the input data to the network. During training, the weights are adjusted so that each output unit represents a set of related inputs. The next subsection will describe the most commonly applied variant of supervised training.

2.2.1 Backpropagation Training

Of the many variations of neural networks, the backpropagation trained network is the most widely applied. Applications of backpropagation include text to speech conversion [7], sonar echo classification [8], cloud classification [9], and ionospheric radar return classification [10], and satellite image processing [11].

Backpropagation networks are hierarchical networks consisting of several layers of independent neurons. These layers are designated as: the input layer, one or more hidden layers, and the output layer. Figure 3 on page 7 illustrates this architecture using a three layer network. As described in [12–14] the output of a single neuron (except, of course, the input neurons) is given by

$$o = f(\sum_{i=0}^{N-1} w_i x_i)$$

where x_i is the value for input i, w_i is the corresponding weight, and f is a monotonically increasing, nonlinear, differentiable function, traditionally the sigmoid function

$$f(a) = \frac{1}{1 + e^{-a}}$$

Training is accomplished by applying the following algorithm until the mean squared error is less than some specified threshold, that is, until the quantity

$$\sum_p \sum_k (tar_{pk} - o_{pk})^2$$

is minimized, where p ranges over the input patterns, k over the output units, tar represents the target response, and o is the actual response:

Backpropagation Training Algorithm

1. Initialize the weights to small random values.
2. Present the inputs and target outputs.
3. Calculate output values for all of the units, beginning with the input units and progressing through the network.
4. Starting at the output units and working back through the network adjust the weights from unit j to unit k by

   $$w_{kj}(t+1) = w_{kj}(t) + \eta \delta_k o_j$$

   where η is the learning rate and

   $$\delta_k = o_k(1 - o_k)(tar_k - o_k)$$

   for output units and

   $$\delta_j = o_j(1 - o_j) \sum_k \delta_k w_{kj}$$

   for units in the hidden layer.
5. Repeat steps 2-4 until the error criteria given above is met.

One of the problems which is commonly used as a simple demonstration of backpropagation is the encoder illustrated in Figure 3, the output of the PlaNet [15] simulator under X-Windows. The goal in this example is to reproduce the input values at the outputs through the bottleneck of two hidden units. This implies that the hidden units will have to evolve some internal representation which encodes the input values. Figure 4 shows the same network after training. Note that the hidden units have formed a binary encoding of the inputs, and that the output units properly track the corresponding inputs.

Since backpropagation is one of the more mature neural network training mechanisms, several investigators have compared the effectiveness of backpropagation with other pattern recognition techniques [16–18]. The sense of these studies is that backpropagation is generally competitive with statistical techniques. Strengths include the absence of a priori distribution assumptions and the ability to interpolate ("generalize") in the presence of deviations from the training set [16].

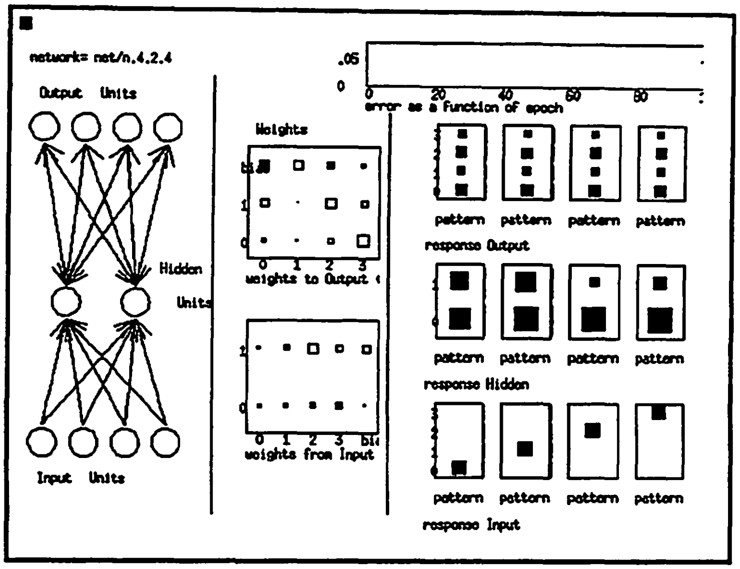

Figure 3: PlaNet 4.2.4 Encoder Before Training

2.3 "Biologically Plausible" Networks

While the neural network model described above is only metaphorically related to actual neural activity, other models have been developed which retain a more faithful likeness to their biological counterparts. Examples include Yao and Freeman [19] who simulated pattern recognizers based upon chaotic interactions between oscillators modeled after the neural interactions within the olfactory system, Fukushima et.al [20] with a simulation of human visual recognition, and Beer [21] with a study of a simulated cockroach.

In Beer's study of computational neuroethology [21], he describes a "hard-wired" neural network which models several behaviors of the American cockroach. Know as "Periplaneta computatrix", or computer cockroach, this simulated animal can walk, explore, and seek food. While it uses a computer for simulation, none of these behaviors require a classical digital computer, the ability to perform these actions is distributed in the interneural connections and excitability.

More recently, Beer has applied aspects of evolutionary computing to allow P. Computrix to evolve some of these behaviors. Evolutionary computation is the subject of the next part of this essay.

2.4 Evolutionary Computing

Evolutionary computing attempts to abstract the salient features of the process which, presumably, produced the intelligent creatures found on earth today. The following items, abstracted from Atmar [22], form some basic assumptions:

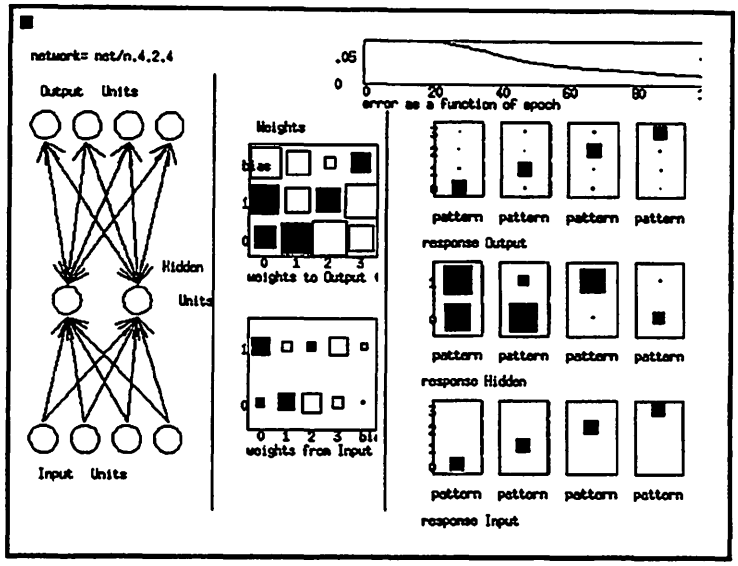

Figure 4: PlaNet 4.2.4 Encoder After Training

- Darwinian evolution is an optimization algorithm. It attempts to find the set of organisms which can survive in a given environment.
- Reproduction is only a approximate duplication due to thermodynamic errors inherit at temperatures above absolute zero.
- Selection acts to statistically cull the least fit individuals from the population

The operation of evolution can be viewed as a search for the optimal population satisfying some criteria of fitness, assumed to be the capacity to survive in biological systems. In computational intelligence these features are used to optimize searching, evolve connections and strengths for neural networks, and evolve programs themselves. In this section I will summarize two of the currently popular evolution-based algorithms in the current literature.

2.4.1 Genetic Algorithms

Genetic Algorithms are abstracted from the mechanisms of genetics information transfer in nature. This includes selection, evaluation, mutation, and sexual reproduction.

The basic GA models the chromosome as a bit-string, with each bit in the string representing a metaphorical gene. The problem of interest is encoded onto this chromosomal representation and the GA itself applied as follows:

1. Randomly initialize a population of chromosomes
2. Test each chromosome string to determine fitness

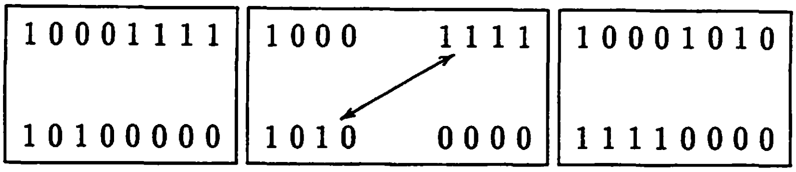

(a) Parents (b) Crossover (c) Children

Figure 5: Crossover

3. Select parent chromosomes from the population in proportion to fitness
4. Reproduce (perhaps with mutation)
5. Repeat from 2 until chromosome produces desired result to within a given tolerance

During the reproduction stage "genetic" information is modified in, typically, in one of two ways, mutation and sexual transmission. Mutation simply inverts a single randomly chosen bit in the chromosome string. This insures some diversity independent of the population. Sexual reproduction involves the selection of two candidate "parents" in proportion to fitness, randomly chooses a point in the string, and swaps the substrings formed by breaking the chromosome at that point as illustrated in Fig. 2.4.1.

While the exact nature of the search is not well understood, practical applications have emerged. Goldberg [23] describes applications such as General Electric's turbine design optimization, leading to more efficient engine performance on Boeing's 777. New Mexico State University has developed "Faceprint" [24], which subsitiutes for the traditional police artist. By using the witness as the fitness function, increasingly accurate pictures of a suspect can be developed by the genetic search. Other applications include optimizing power distribution, target detection, and communication network design.

2.4.2 Genetic Programming

Genetic Programming (GP) is an extension of the genetic algorithm which actually evolves programs which solve particular problems. Invented by John Koza [25] it uses program operators and constants instead of bit strings as metaphorical chromosomes. The basic algorithm is similar the GA as follows:

1. Randomly generate a population of programs
2. Run each of these programs and evaluate their results with respect to the desired outcome (fitness)
3. Reproduce the program with mutation and/or crossover

4. Repeat until desired result is reached to within tolerance

Mutation simply changes one (or more) of the operators or constants to some other randomly chosen value. The crossover operator consists of taking randomly chosen subprograms and swapping them forming new programs. This is illustrated in Fig. 2.4.2 where the program fragments along the arrows marked 'X' are swapped.

Applications of Genetic Programming include many of those in genetic algorithms, including aesthetic art generation [26, 27], learning to balance an inverted pendulum ("broom balancing"), and automatic image target recognition [28].

2.5 Artificial Life

Artificial Life research attempts to abstract the characteristics of life, and reproduce it in some computational form. Farmer et al. [29] describes some these attributes:

- Life is a pattern in spacetime (e.g. most of our cells are replaced in our lifetime).
- Self-reproduction
- Information Storage of self-representation (e.g. DNA)
- Metabolism
- Able to interact with environment
- Interdependence of parts forming the organism
- Stable under perturbations and small changes
- The ability of the lineage to evolve.

Of course, like AI, Artificial Life defies definition. Or, as Langton asserts [30]:

> Although AI has not yet achieved anything that even its most ardent supporters would call genuine machine intelligence, AI has completely changed the way in which scientists think about what it is to be "intelligent", and has, therefore, made a major scientific contribution, even though it hasn't achieved its overall goal.
>
> Similarly, Artificial Life will force us to rethink what it is to be "alive."

Those involved in research in the Alife community recognize two different claims, known as the strong claim and the weak claim. The weak claim asserts that anything produced is a simulation which may help to explain certain properties of life. The strong claim asserts that the computer programs will eventually gain the state of being actually alive. That is, if you abstract all of the relevant properties of what we call biological life, encapsulate them in a machine, then the machine is alive.

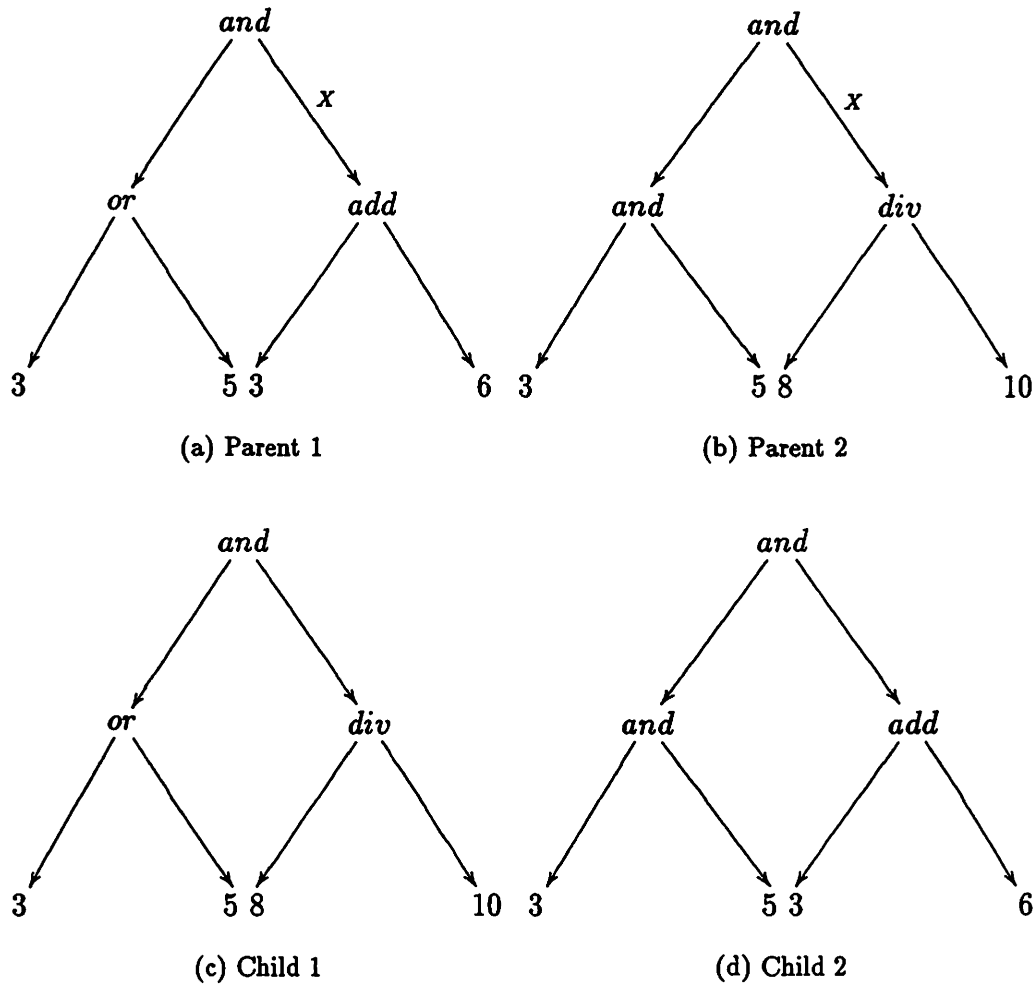

(a) Parent 1

(b) Parent 2

(c) Child 1

(d) Child 2

Figure 6: Crossover

3 Social Implications

The preceding section describes the technology driving the current research in Computational Intelligence. In this section I will try to give an impression of potential consequences. I believe that these consequences are possible *independent* of actually creating artificial intelligence or artificial life.

Farmer [31] sets the stage for the social analysis of these technologies as follows:

> With the advent of artificial life, we may be the first species to create its own successors. What will these successors be like? If we fail in our task as creators, they may indeed be cold and malevolent. However, if we succeed, they may be glorious, enlightened creatures that far surpass us in their intelligence and wisdom. It is quite possible that, when the conscious beings of the future look back on this era, we will be most noteworthy not in and of ourselves but rather for what we gave rise to. Artificial life is potentially the most beautiful creation of humanity. To shun artificial life without deeper consideration reflects a shallow anthropocentrism.

Claudill [32] describing the possibility (inevitability ?) of intelligent, cognizant, robots raises similar issues:

> Will androids become too humanlike? Will they become prey to the evils in mankind as well as our virtues? Is an android likely to develop prejudices, become selfish, or turn violent?

In the event the android develops emotions such as love and hate she asks if you would want your son or daughter to marry a sexually functional robot? And finally, "if we are the creators of this race of independent intelligent beings, does that mean that we are on par with God? What is the difference between breathing life into a lump of clay to make a human being, and constructing an intelligent, self-aware android?".

Nobel Laureate, Francis Crik, reinforces this potential in his book *The Astonishing Hypothesis,* [33]. In the words of Crik,

> The Astonishing Hypothesis is that "you", your joys and your sorrows, your memories and your ambitions, your sense of personal identity and free will, are in fact no more than the behavior of a vast assembly of nerve cells and their associated molecules.

and then describes the anticipated role of traditional religions as:

> History has shown that mysteries which the churches thought only they could explain (e.g., the age of the earth) have yielded to a concerted scientific attack. Moreover, the true answers are usually far from those of conventional religions. If revealed religions have revealed anything it is that they are usually wrong. The case for a scientific attack on the problem of consciousness is extremely strong. The only doubts are how to go about it, and when. What I am urging is that we should pursue it now.

Indeed, the pursuit is in full force. Prominent topics of discussion at the recent World Congress on Neural Networks spent a considerable amount of time wrestling with the ethical treatment of a conscience machine. The assumption is that it is simply a matter of time. Already sensors are being tested which are based on an array of neural cells grown on a silicon substrate (citation to be provided in final paper). Experiments are also in progress to provide direct neural control of prosthetic devices, blurring the distinction between man and machine.

This last point is important. Mazlish [34] argues that man and machine are co-evolving in symbiosis. He identifies three areas where anthropocentrism has been challenged, then adds his own. The first, Copernicus, put man in his proper place in the universe – as a small fragment in a vast domain. The second, Charles Darwin, erased the unique place man had assumed in creation by placing him simply as the descendant of other animals. The third, Sigmund Freud, who with psychoanalysis showed that "the ego ... is not even master of its own house". And finally, the fourth, that "the sharp discontinuity between humans and machines is no longer tenable, in spite of the shock to our egos". Once this discontinuity between man and machine is removed, the way is paved for the machine to equal or surpass the human on the evolutionary ladder.

It must be noted that, as stated at the onset of this discussion, there is disagreement over the limits of Artificial Intelligence. It is by no means apparent that these goals will be achieved. The hope is that by modeling nature's implementation of intelligence computational intelligence will result. I believe that, even if computational intelligence is totally discredited in the future the basic assumptions as illustrated above can have consequences. Even if, for example, that each of the discontinuities that Mazlish mentions proves inaccurate, their assumption could, in the limit, reduce mankind to a level below the machine. And if machines should become conscious what would be our ethical response to them?

Another potentially serious consequence, again in the limit, is to consider anyone who challenges the ethics of deploying intelligent machines an impediment to the natural order of things, relatively unfit, and therefore expendable. If we believe that the machine is our successor, and we have no solid foundation for our ethics, then we can view the individual dissenter as little more than an illness which must be inoculated against. In principle this is little different than the attitudes preceding the holocaust of World War II.

4 The Classroom

The preceding sections outline the environment in which computer science students must develop their outlook on life and their discipline. This very paper is largely the result of questions asked by students in the classroom. Typical questions include:

- Is machine intelligence something that God would want us to study?

- Can a machine really be intelligent? If so, how should we treat it?

- If we are just like a machine, what about our free will?
- How can we ensure ethical treatment of people if machines become intelligent?
- If evolution works in abstract form, why can't it work in nature?

In some sense all of these questions are less about machines than they are about our place in a rapidly changing world. This is like searching for a stable rock to protect from a raging storm, a framework from which to operate. It is my belief that the bible can provide a stable platform for exploring these questions, but that some of the answers may require a face-to-face discussion with God, perhaps on the new earth.

The question of free will is a difficult one. If you work as a deterministic machine how can you have free will? Assuming that you have free will at all, does it exist in any domain or is it restricted? There is evidence that the brain works within the realm of chaos, could this be involved? I don't know. What I can say with confidence is that we are assured the ability to choose where our salvation is concerned. It is not clear that scripture supports more than that. If I understand, for example, what Paul is saying in Romans 6:16 then I am either a slave to sin or a slave for obedience under grace apparently restricting the domain of free will.

Concerning questions involving God's potential concern about the study of machine intelligence, or the question of the possibility of a machine becoming intelligent one is tempted to glibly quote truisms such as "all truth is God's truth" [35] and let it go at that. This would, I believe, be a disservice to our students, for I believe that the real question being asked is: "if a machine is actually intelligent does that make me less of a human being?".

I am not sure that this identification of self with intelligence is an overtly Christian view. It is, however, pragmatic. Society at large certainly values intelligence as a defining attribute. Our distance from other creatures is often cited as our intelligence and self-awareness. And our ability to cloth and feed our families is to some extent dependent upon intelligence in a modern culture. Yet, upon reflection, this strikes me as a subtle, and unconscious, form of idolatry where the intellect is held supreme. Instead, as Christians, I would suggest that our intellect is a gift from God, to be used in his service. It is not the defining attribute of our humanity. God defined us, he created us in his image, he died for us, and he will return for us. It is in his unimaginable sacrifice for us that we gain our identity, and it is for this reason that I believe we have nothing to fear if a machine should become intellectually superior, as unlikely as that may be.

The question of the ethical treatment of people in an environment which considers the intellect the defining human attribute is serious. History has not allowed the Christian a place of esteem in the treatment of people. Likewise, secular society has little compassion for those who would act in the critical role of the prophet. Yet, it seems to me, the Christian comfortable in his relationship with God is in a unique position. If people are important to God, they should be important to the Christian. For the Christian the value of the individual is not based on current social context, it is based on God's command to love each other.

5 Conclusion

The understanding of our role as God's creatures and his sacrifice for us affords the Christian a stable platform, supported by scripture and informed by the Holy Spirit, that can, I believe, bear the weight of radical change. People count to me, as a Christian, because they are important to my God. To the extent that technology affirms life, as Jesus affirmed life, it is good. This does not mean that specific answers to specific questions will be unambiguously apparent. It does not mean that uniformity without dissent will emerge. It does mean that there is a place to stand from which to approach these difficult questions.

References

[1] H. Moravec, "IEEE spectrum book review," *IEEE Spectrum*, vol. 31, Sept. 1994.

[2] G. Stock, *Metaman: The Merging of Humans and Machines into a Global Superorganism.* Simon and Schuster, 1993.

[3] R. L. Paden, "Artificial intelligence: From the foundations of mathematics to intelligent computers," in *Proceedings of the Institute for Christian Teaching: Christ in the Classroom, 1992*, Institute for Christian Teaching, 1992.

[4] S. Forrest, "Emergent computation: Self-organizing, collective, and cooperative phenomena in natural and artificial computing networks," in *Emergent Computation* (S. Forrest, ed.), (Cambridge, MA), MIT Press, 1991.

[5] B. Muller and J. Reinhardt, *Neural Networks.* New York: Springer-Verlag, 1990.

[6] B. Kosko, "Unsupervised learning in noise," *IEEE Transactions on Neural Networks*, vol. 1, pp. 44–57, March 1990.

[7] T. J. Sejnowski and C. R. Rosenberg, "Nettalk: A parallel network that learns to read aloud," Tech. Rep. JHU/EECS-86/01, The Johns Hopkins University Electrical Engineering and Computer Science Report, 1986.

[8] P. R. Gorman and T. J. Sejnowski, "Learned classification of sonar targets using a massively parallel network," *IEEE Transactions on Acoustics,Speech, and Signal Processing*, vol. 36, July 1988.

[9] J. Lee, R. C. Weger, S. K. Sengupta, and R. M. Welch, "A neural network approach to cloud classification," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 28, September 1990.

[10] V. G. Sigillito, S. P. Wing, L. B. Hutton, and K. B. Baker, "Classification of radar returns from the ionosphere using neural networks," *Johns Hopkins APL Technical Digest*, vol. 10, no. 3, 1989.

[11] J. Wolfer, J. Roberg'e, and T. Grace, "Robust multispectral road classification in landsat thematic mapper imagery," *Proceedings of the World Congress on Neural Networks*, jun 1994.

[12] R. P. Lippmann, "An introduction to computing with neural nets," *IEEE Acoustics Speech and Signal Processing Magazine*, vol. 3, no. 4, 1987.

[13] D. E. Rumelhart and J. L. McClelland, *Parallel Distributed Processing*, vol. 1. Cambridge, MA: MIT Press, 1986.

[14] J. Hertz, A. Krogh, and R. G. Palmer, *Introduction to the Theory of Neural Computations.* Redwood City, CA: Addison Wesley, 1991.

[15] Y. Miyata, *A User's Guide to PlaNet Version 5.6 A Tool for Constructing, Running, and Looking into a PDP Network.* Yoshiro Miyata, 1990. PlaNet distributed by miyata@boulder.colorado.edu.

[16] J. A. Benediktsson, P. H. Swain, and O. K. Ersoy, "Neural network approaches versus statistical methods in classification of multisource remote sensing data," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 28, July 1990.

[17] J. Key, J. Maslanik, and A. Schweiger, "Neural network vs. maximum likelihood classifications of spectral and textural features in visible, thermal, and passive microwave data," in *International Geoscience and Remote Sensing Symposium*, (College Park, MD), p. 1277, 1990.

[18] P. D. Heermann and N. Khanzenie, "Application of neural networks for classification of multi-source multi-spectral remote sensing data," in *International Geoscience and Remote Sensing Symposium*, (College Park, MD), p. 1273, 1990.

[19] Y. Yao and W. Freeman, "Pattern recognition in olfactory systems: Modeling and simulation," in *Proceedings of the International Joint Conference on Neural Networks*, 1989.

[20] K. Fukushima, S. Miyake, and T. Ito, "Neocognitron: a neural network model for a mechanism of visual pattern recognition," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. SMC-13, 1983.

[21] R. D. Beer, *Intelligence as Adaptive Behavior An Experiment in Computational Neuroethology.* Academic Press, 1990.

[22] W. Atmar, "Note on the simulation of evolution," *IEEE Transactions on Neural Networks*, vol. 5, Jan. 1994.

[23] D. E. Goldberg, "Genetic and evolutionary algorithms come of age," *Communications of the ACM*, vol. 37, no. 3, 1994.

[24] C. Caldwell and V. Johnston, "Tracking a criminal suspect through face-space with a genetic algorithm," in *Proceedings of the fourth International Conference on Genetic Algorithms*, 1991.

[25] J. Koza, "Genetic evolution and co-evolution of computer programs," in *Artificial Life II*, Addison Wesley, 1992.

[26] K. Sims, "Artificial evolution for computer graphics," *Computer Graphics*, vol. 25, July 1991.

[27] S. Todd and W. Latham, *Evolutionary Art and Computers.* Academic Press, 1992.

[28] W. A. Tackett, *Recombination, Selection, and the Genetic Construction of Computer Programs.* PhD thesis, Universtiy of Southern California, 1994.

[29] J. D. Farmer and A. Belin, "Artificial life: The coming evolution," in *Artificial Life II*, Addison Wesley, 1992.

[30] C. G. Langton, "Introduction to the artificial life II proceedings," in *Artificial Life II*, Addison Wesley, 1992.

[31] D. Farmer, "Untitled." 1992.

[32] M. Caudill, *In Our Own Image: Building an Artificial Person.* Oxford, 1991.

[33] F. Crick, *The Astonishing Hypothesis: The scientific search for the soul.* New York: Scribners, 1994.

[34] B. Mazlish, *The Fourth Discontinuity: the Co-Evolution of Humans and Machines.* Yale University Press, 1993.

[35] A. F. Holmes, *The Idea of a Christian College.* Wm. B. Eerdmans, 1975.