

02D021I

Project Number: 46-AMF-A101

02D021I

MaiNSTREAM: Artificial Intelligence and Society

An Interactive Qualifying Project Report

submitted to the Faculty

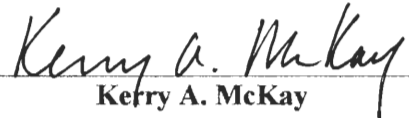
of the


WORCESTER POLYTECHNIC INSTITUTE

in partial fulfillment of the requirements for the

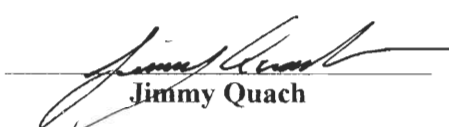
Degree of Bachelor of Science

by


Kerry A. McKay


Robert Sartorelli


Lindsey Perullo


Jimmy Quach

Date: April 15, 2002

Approved:


Professor Adina Florea, Advisor

- 1. Artificial Intelligence
- 2. Society



Artificial Intelligence & Society

An IQP by: **Kerry McKay**
Robert Sartorelli
Lindsey Perullo
Jimmy Quach
Advised by: **Adina Florea**



Abstract

The goal of our project is to study the impact of artificial intelligence on current society and to make some predictions on how the field will affect society in the future. Through a brief look at the background and origins of the field, we will assess where AI has been. We will examine AI's current uses and its contemporary influence in the media. Also, we have conducted a short survey, through it we will be able to show how AI is viewed by people now. Lastly, we will discuss some moral and ethical issues pertinent to the AI field and its future along with some risks and perils it could present. We predict that artificial intelligence will gain in importance as a field of study, becoming intertwined with almost every other aspect of life until we are nearly reliant on machines for our day-to-day lives. Despite this somewhat scary thought, we believe that ultimately AI will provide society with a means to improve the quality of life.

Table of Contents

Introduction	4
1. A Brief History of AI	6
1.1 When did it become reality?	6
1.2 How did “Artificial Intelligence” come about?	6
1.3 Early examples.....	7
1.4 Summary	8
2. Alan Turing – Can machines think?	9
2.2 A brief look at some objections	10
2.3 The imitation game	11
2.4 Can a machine learn?	13
2.5 Summary	14
3. Approaches	15
3.1 Learning in Machines	15
3.2 Neural Networks	16
3.3 Simulation vs. Emulation.....	18
3.4 Summary	19
4. AI Use Today	20
4.1 Games	20
4.2 Toys.....	23

4.3 Tools in medicine, military, and music.....	28
4.4 Summary.....	30
5. Candidates for the Turing Test	31
5.1 Hal.....	31
5.2 Alice.....	34
5.3 Ella	36
5.4 Summary	38
6. AI in the Media.....	40
6.1 Movies.....	40
6.2 Television.....	52
6.3 Literature.....	57
6.4 Summary	63
7. Public Opinion.....	65
7.1 Our Predictions	65
7.2 Survey Results	69
7.3 Survey Interpretations.....	70
7.4 Summary	73
8. Moral and Ethical Questions	75
8.1 Definitions and Goals	75
8.2 Philosophical Questions.....	79
8.3 Social Aspects.....	85
8.4 Responsibility	87

8.5 Summary90

9. Risks and Perils of AI 92

9.1 Reliability.....92

9.2 Human-Computer Compatibility95

9.3 Is AI a Threat?97

9.4 Summary98

10. Conclusions..... 100

Appendix A – Methodology..... 103

Appendix B – Survey 105

Appendix C - Results from the survey..... 109

Appendix D - Comments 119

Appendix E – Ella Log..... 123

Appendix F – Web Screenshots 127

Bibliography 130

Introduction

In the twenty-first century, we are surrounded by technology that is progressively getting “smarter”. It’s becoming more advanced, more automated, and more complex. But how smart will technology get? Everyone has heard of artificial intelligence (AI). It has been a predominant part of science fiction for many years, but is it possible?

Not too many people know the origins of this concept. In this paper, we will discuss the history behind artificial intelligence, as well as take a look into Turing’s influential paper “Computing Machinery and Intelligence”.

One of the major problems in artificial intelligence is “how do we do it?” “Approaches” will discuss the two main ways that this problem is being tackled, and also compare the two from the views of simulation and emulation.

Artificial intelligence is all around us in our everyday lives, as many people are finding out. From toy dogs to surgeons, robots are becoming more and more mainstream, causing some people to feel uneasy about the future of humanity. Some people are afraid of a robot-takeover, such as the ones seen in the films *The Matrix* and *Terminator*. “AI Use Today” will explain where it’s being used, and what it’s being used for. The chapter on “Public Opinion” will show the public’s view as gathered by a survey of eleven questions. The possible negative effects of artificial intelligence in our society will be examined in “Risks and Perils of AI”.

There is an annual competition sponsored by Hugh Loebner that follows Turing’s imitation game, also known as the Turing test. Two of the seven 2001 contestants are discussed in “Candidates for the Turing Test”, along with another creation that will be perfectly suited for the test when it “grows up”.

Introducing artificial beings to human society opens a Pandora's box of ethical and moral concerns. Should they be treated as equals or lesser beings? Will we be responsible for their upkeep, or are they on their own after creation? Is artificial intelligence just a quest of mad scientists? Questions such as these are very difficult to answer. These issues, among others will be discussed in "Moral and Ethical Questions".

The aim of this project is to learn about the interaction between human society and artificial intelligence and make predictions for the future of the two existing together. We hope to show that artificial intelligence is not a horrifying concept despite what some of the stories say. In industry and business, from our cars to our computers, and everything in between; artificial intelligence will be employed to help us. It will become intertwined with almost every other aspect of life until we are nearly reliant on machines for our day-to-day lives. Despite this somewhat scary thought, we believe that ultimately AI will provide society with a means to improve the quality of life.

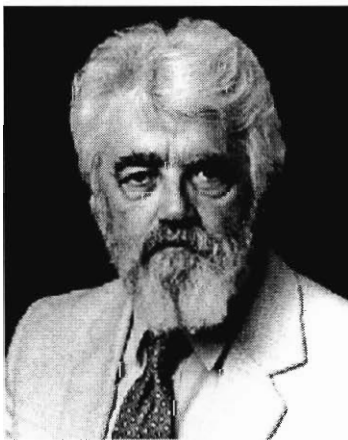
1. A Brief History of AI

The idea of artificial intelligence goes as far back as the greek myths. Myths of Hephaestus, the god of fire and the patron of craftsmen, show proof of the idea of intelligent robots. It has since then evolved into something much more. The concept of thinking machines has been around for many many years, but “artificial intelligence” is an idea formed in the last century.

1.1 When did it become reality?

It was around 1945 that artificial intelligence in its symbolic aspect became a specific field of study. Within the following decade, developments like Norbert Wiener’s mathematical theory of feedback in biological and engineered systems, which created the field of cybernetics, helped broaden the study of artificial intelligence and clarify the concept that intelligence is the process of receiving and processing information to achieve goals.

1.2 How did “Artificial Intelligence” come about?



The term “artificial intelligence” wasn’t used to describe this field of study until 1956, when John McCarthy used it as the topic name of the Dartmouth Conference, which was the first conference devoted to the concept which we now refer to as “artificial intelligence”.

Figure 1.1.1 – John McCarthy

1.3 Early examples

One of the first “intelligent” machines designed (but never completed) was Charles Babbage’s Analytical Engine. Proposed in 1832, it was to be a massive brass, steam-powered mechanical computer used to execute not only arithmetical calculations, but all those of analysis, if their laws are known. It was to be an automatic general



Figure 1.2.1 – Charles Babbage

purpose machine that could be programmed by the user to execute instructions in any required order. Babbage continued to modify and refine his design until his death in

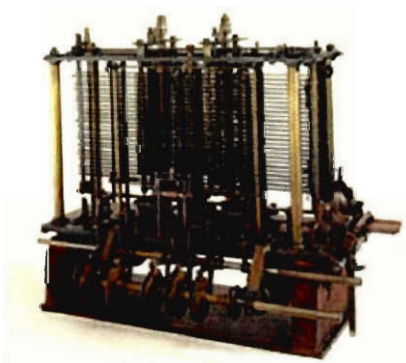


Figure 1.2.2 – Portion of the Mill of Analytical Engine

tolerance with the technology at that time.

1871. By then the machine consisted of four main parts: the store (memory), the mill (computational unit), the punched card reader (input section), and a card puncher (output section). Babbage had all the essential ideas for a modern computer but wasn’t able to produce most of the parts in their needed degree of precision and

The first artificial intelligence program, the Logic Theorist, was written by Allen Newell, Herbert Simon, and J.C. Shaw in 1956. Using a combination of searching, goal-oriented behavior, and application rules, it proved mathematical theorems. It was able to recreate many previously proven theorems. In one case, it even derived with a more elegant proof than the one devised by humans. The Logic Theorist was written in the information processing language (IPL), which had pointers between related pieces of information to mimic associative memory.

In 1966, Eliza was created by Joseph Weizenbaum. She was programmed as a therapist that recognizes keywords from the patient and gives a response using the keywords in some manner. Eliza can't plan out actions, learn, have the ability to reason or understand her own motives but surprisingly her interviewees would create strong emotional bonds with her. This left Weizenbaum fearing the possibility of society easily giving artificial intelligence the control of human affairs.

1.4 Summary

Artificial intelligence began as a mythical idea and has progressed into a practical reality. With the help of creative and brilliant minds in designing and putting intelligence into machines, artificial intelligence has quickly advanced and infiltrated the mainstream. These accomplishments are only the beginning of a vast and incredible journey.

2. Alan Turing – Can machines think?



Figure 2. 1 – A. M. Turing

Alan Turing was a famous mathematician who is regarded as the father of computer science. In 1950, Turing’s article “Computing Machinery and Intelligence” was the first article to bring a great amount of attention to the field of AI and get people thinking about it. This article dared to ask the question “Can machines think?”, which is a controversial question even to this day. His ideas are still important and are

the source of both reference and criticism.

2.1 What defines thinking?

Unfortunately, there is no “clear cut” answer to this question. According to Webster’s Dictionary, to think is “to exercise the powers of judgment, conception, or inference”, and “to call to mind”, as well as other definitions. Dictionary.com defines “think” as “to have or formulate in the mind”, “to decide by reasoning, reflection, or pondering”, or “To judge or regard; look upon”.

The “thinking” done by computers for AI purposes seems to entail processing stimulus data and coming to sort of idea or conclusion. The stimulus could be in the form of text, voice, light, touch, anything that the computer in question can detect. An idea or conclusion doesn’t necessarily have to be reached right away, the computer should be allowed to question, but something should be understood, whether it be a final conclusion, a new idea, or the realization that more information is needed in order to reach an idea or conclusion.

2.2 A brief look at some objections

Not everyone believes that a thinking machine is possible, and for many different reasons. Some people have logical, well thought out reasons, some have completely illogical reasons that have nothing to with the technology involved, and some people think it's impossible simply because they want it to be. In his article, Turing wrote a section on objections such as these. Two of the logical ones will be discussed here.

The mathematical objection is based on Gödel's famous Incompleteness theorem. In it, he showed that "in any sufficiently powerful logical system statements can be formulated which can neither be proved nor disproved within the system, unless possibly the system itself is inconsistent" (Turing, 1950).



Figure 2.2.1 – Kurt Gödel

Another well-reasoned view is the objection that relates to the nervous system. Computers are discrete-state machines, while the nervous system is continuous. A discrete system can not accurately imitate a continuous system because even the smallest error can greatly change the result. Neurons fire impulses from one to another in order to make us think. If just one of the neurons receives an impulse that is off, the entire impulse chain is off. Once this happens, the result dramatically changes; therefore, a discrete-state machine fails to simulate a continuous nervous system.

2.3 The imitation game

In Turing's article, he describes a game called the "imitation game". It was designed to determine which player is real, and which is the imitator. There are three participants in this game: the interrogator, the person, and the computer. The person and computer are in a different room than the interrogator. The interrogator asks one of the other players a question, and they respond via text to the interrogator. After a series of questions, the interrogator must decide who is the computer, and who is the person. If the computer fools the interrogator into thinking it's the person, it passes the test. This is very difficult, for the interrogator may ask any question. Human knowledge is quite vast and diverse. Any computer should be capable of "number crunching", but what about creating new ideas, or giving thoughts on literature? If a machine were asked, "What do you think Shakespeare was saying in 'Romeo and Juliet'?", how would it respond? A.L.I.C.E said "I haven't thought much about Shakespeare was saying in Romeo and Juliet, so I don't have a good answer for you." While this is not necessarily a bad response (other than the grammatical error), but it is a bit of an easy way out. Most people, at least in America, have read the play, seen it performed on stage or on film, or at least know the main idea of the story. They would then take a moment to think of the plot and come up with some answer, or simply say "I don't know".

What Turing called "the imitation game" grew to be much more. Today, it is known as the "Turing test" – a test that decides whether something is or is not intelligent. There are people who believe that Turing's paper was misinterpreted; it may have been used in a way that Turing had not intended. Turing never referred to the imitation game

as anything more than a simple game; yet despite this, it's become the Holy Grail for



Figure 2.3.1 – The Loebner Prize gold medal

some AI developers.

In 1990, Hugh Loebner began a competition in which the candidates are subjected to the Turing test. The award is given to the creator(s) of the candidate that most resembles a human being. The first person to create a robot that can convince half the judges that it is a human through speech will receive \$100,000 and a gold medal (Figure 2.3). A silver medal and \$25,000 goes to whoever accomplishes the same with text instead of voice. Finally, bronze is awarded every competition to whoever creates the most human-like robot. The competition will end once someone receives the gold.

But when does something go from having the ability to think to being intelligent? Some people find this inability to distinguish the computer from a human as enough to classify a machine as intelligent. Others, however, feel that intelligence is not achieved until the machine is self-sufficient; not needing humans to maintain it, but being an entity which is like an animal built of circuits and metal opposed to the organic make-up of all earth's creatures.

Fooling humans may not be enough to classify something as intelligent. As Whitby points out in his 1997 paper "Why the Turing Test is AI's Biggest Blind Alley",

we know very little about human intelligence. In any other field of study, would we model something that we don't know about? A very good example of this was given, comparing the natural flight of a bird to the flight of airplanes. Research in natural flight was very unhelpful in the quest for artificial flight, but understanding the principles of aerodynamics led to the creation of working aircraft.

The slugbot (<http://www.ias.uwe.ac.uk/goto.html?slugbot>) is a good example of a self-sufficient machine. Structurally, it's comprised of a body, wheels for mobility, a basket, and an arm mounted on a turntable with a claw at the end. It utilizes a red light on the end of the arm to spot slugs, which are easily recognizable under red illumination with a CMOS image sensor. When it finds a slug, it uses the claw to pick it up, and then deposits it in the basket. The slugbot returns to a "charging station", which it may share with other slugbots, where the slugs are turned into a form of energy that the slugbots can use. This form of energy conversion allows it to operate without humans supplying it with electricity.

2.4 Can a machine learn?

Many people believe who are not in the technical field or do not know much about AI feel that machines are incapable of learning as a human would. Many researchers familiar with the field however are convinced that not only is it possible to model the human learning process in machines but that it is inevitable.

There are three basic components of the learning machine. The first is the initial state of mind, which is similar to the birth of a human being. The mind is empty and completely void of what we'd call knowledge. Education is the next component. Just the

same as a human or other creature, a computer must be taught information to use later on in life. The final part is experience. Even if we were told something is true, we don't always believe it. Experience is somewhat of a proof in this sense. Only through experience can we really learn what we are taught.

Turing discussed the advantages and disadvantages of imitating the mind of a child vs. the mind of an adult using this idea. A child learns more through experience. Imitating an adult mind would have information as "fact", where as it is more of a guideline in the mind of a child. Once educated and subject to experiences, the child-like computer will evolve into an adult-like one.

2.5 Summary

Alan Turing was a major figure in the study of artificial intelligence. His 1950 paper "Computing Machinery and Intelligence" has had quite an impact on research, and some believe that his imitation game is the model to follow on the way to achieving AI. This belief has spawned into an annual competition in which robots and their creators compete for money, a medal, and a title. While the lines that separate thinking from not thinking and intelligent from unintelligent are still quite gray, computers are definitely on their way to reaching a new level of processing.

3. Approaches

In most problem-solving processes, not everyone attacks the issue in the same way. There are many approaches, and often many solutions to the same problem. In the case of artificial intelligence, the two main methods of attack are emulating the human mind through a series of algorithms, and simulating processes of the human mind and nervous system.

3.1 Learning in Machines

Machine learning is defined as the change in a system that is adaptive. This enables the system to do the tasks more efficiently the next time they come about. In an overview by a man named Patrick Doyle, he goes into the “four main components” of machine learning. The first aspect of learning is called the learning element, which is the component of the agent that is responsible for developing performance. The second is the performance element, being the component that chooses which actions it is to take. Next is the critic. This is what tells the learning element what and how the agent is doing. Finally, there is the problem generator. Designed to improve the element, the problem generator suggests what actions (causes) can lead to what effects.

Doyle also mentioned that there are four issues to consider in the techniques of designing a learning machine. There are the components, the parts of the performance element that need to be improved. There is the representation of those components. Another important issue is the feedback that becomes available to the learning machine. Also, the prior information is an issue because it is important to know what the machine already knows.

There are various types of learning, one being “speed-up” learning. This involves data caching, which requires no additional input. Yet, it does improve the agent’s performance and “knowledge” over some time. Another type is learning by taking advice. This is where the system or learning machine is able to reason the new information being added to its already existing prior knowledge. Learning from examples is similar. The learning machine learns concepts for certain instances, and adds them to its existing knowledge. An additional inductive way of learning is clustering. This is where the system classifies “natural classes” from data instances. Analogy is machine learning by way of knowledge being transferred to the system to some domain from a database. Finally, there is learning through discovery. It is a method of gaining knowledge through inductive as well as deductive learning, where the agents learns somewhat on its own. The learning process is deductive by proving theorems.

3.2 Neural Networks

Another approach to the artificial intelligence problem is to simulate the cells of the human brain interacting. This breed of engineering is called neuromorphic engineering. Neural networks are not only used for artificial intelligence purposes. Today, they are commonly used in areas such as system modeling, pattern matching, and data compression.

They differ from conventional learning machines in the sense that they do not use a set of algorithms to solve problems. Instead of storing information in memory like the average computer, these neuromorphic creations store them in a similar manner that we

humans do; store it in a collection of neurons and synapses. The neuron structure has four key elements:

- Cell body (with nucleus)
- Axon
- Dendrites
- Synapse

The main focuses for this discussion are the last three bullets. The neuron becomes “excited” when an electrical impulse is received by the dendrites. If the signal is significantly greater than a certain threshold, then the dendrites then send an electrical spike down through the neuron to the axon. The signal is then converted back to electrical signal that is sent over the synapse to be received by other neurons.

When a human learns something new, the weights of the synapses between the affected neurons change so that the influence those neurons have on the surrounding neurons changes. The artificial neuron model works the same way. It is comprised of structures that correspond to the ones mentioned above for the human neuron, thus simulating the human brain.

The thought process is a fairly simple concept, but not simple to carry out. The network can be viewed as a collection of nodes (axons) and connections (synapses). Each path has a total weight of the sum of all the connecting segments crossed. The output by each node is calculated by multiply the input by the value of the path taken and then comparing it to a certain threshold determined by a transfer function.

Robots created with this technique can do things that a learning machine on a conventional computer cannot. For example, they can walk and avoid obstacles. This is a

huge step because they can interpret information from the world around them. This area is sure to develop interesting results in the future.

3.3 Simulation vs. Emulation

Two very important paradigms in the implementation of artificial intelligence are simulation and emulation. The two methods share several similarities but also contain some drastic differences. First lets start off with some definitions. Simulation is defined as a “representation of the operation or features of one process or system through the use of another.” Conversely, Emulation is “Effort or ambition to equal or surpass another. Imitation of another.” Applied to computers, a simulation is helpful to analyze a system or a process and make predictions of it whereas emulation is helpful if used as a substitute for replacement of the process it is emulating. - (<http://www.dictionary.com>)

More specific to the realm of artificial intelligence, emulation refers to the attempt to recreate or mimic exactly the process of human thought. This method allows for analysis and prediction of the actual system, in this case, the human brain. Simulation, on the other hand, is not concerned with how to get there, it just deals with results. The path is not important in simulation but the destination is. If human level intelligence or greater is the goal than emulation will attempt to equal or improve upon the system, again, in this case the human brain. The goal of both simulation and emulation as applied to AI is to mimic human thought. In fact, both also strive to improve upon it. Emulation attempts to increase the speed and efficiency of the human brain while keeping the same basic blueprints while simulation looks to create a new process altogether that achieves the same goal in at least an equal fashion. If simulation succeeds it would be an amazing

breakthrough. Human level intelligence would be recreated through a process completely new and dissimilar to our own internal structure. Emulation on the other hand would be the more predictable and perhaps more stable approach.

An example of an approach using the emulation paradigm is a neural network. They attempt to recreate the infrastructure of the human brain with synthetic neurons. The end result is a process very similar to our own biological brains only with much greater speed and accuracy. A counter example to the neural network is the learning machine, which generally utilizes the simulation paradigm. Learning machines follow a process that does not retrace the exact path our own brains make. They use algorithms and heuristics to interpret and learn from the data they are given. This is an attempt to mimic or achieve human level thought through a new and possibly improved method.

3.4 Summary

There is more than one way to try to create an intelligent machine. The human mind could either be simulated or emulated, depending on the desired outcome. If the human mind is to be simulated, then a machine that acts like a human by different processes (a learning machine) is the main method to use. If the mind is to be emulated, then the neural networks approach is better.

4. AI Use Today

4.1 Games

Thus far in the year of 2002, according to an article in WIRED (March 2002), more than a third of the households across America contain some sort of “sophisticated artificial intelligence bot.” Technology is always flourishing, and people from young to old are becoming more computer literate, which is doing wonders for the gaming industry. Many of the video games recently released are very different from those that existed back when games were simple. A games play was as simple as a walk or run through the level, fights with the various “bad guys” of the level, earning points, and defeating the mastermind of it all. Games released in today’s day and age make the player conscious of what mood the character is in, how much stamina they have left to fight, if the character is bored, upset, or happy with its surroundings, etc. This is accomplished by using a variety of techniques and adaptive AI, in some form of self-modifying fuzzy state machines, that reacts to what the game player does. Games like *The Sims*, *Grand Theft Auto III*, *Petz*, and *Black and White* are very good examples of bringing this sort of realness to video games.

In a particular PC video game, *The Sims*, you design virtual buildings and neighborhoods, and then populate them with virtual residents. Each Sim can be created with its own personalized attributes, including skin color, gender, age, sexual preference, and personality traits. From there, they interact with other Sims, based on their world and the personalities that you gave them. They behave like a normal person would; they eat, sleep, play, work, fall in love, and have a family, etc. They learn to interact with

neighbors and other people. They can find jobs by looking in the newspaper classifieds section. If they are unhappy with their current significant other, they will cheat. These Sims will even get bored if there is no entertainment set up for them in their home. They



Figure 4.1.1 – A screenshot of the Sims

either have a negative or positive reaction, and everything that does happen, effects the way the Sim will live its life.

A similar game to *The Sims* is *Petz* or *Monsterz* or any other variation of this PC game. The virtual pet on the user screen acts similar to a real life pet would. You can play with or punish your pet, depending on its behavior, whether it chews up a pair of sneakers, or leaves a mess on your desktop. How the player treats their pet will determine how the pet behaves when it grows up. And you do get to watch it grow and mature. It starts off as a youth, and over time, it will be in its teenage years. Sooner or later, the once adolescent will be a full-grown adult, and so on. If the user treats their pet with love and affection right from the beginning and is not abusive towards the pet, it will grow up to be a calm and cuddly pup. However, if the player were to constantly punish and instigate the pet, then the youth would grow up to be mean and vicious. This is true for any real life pet, which is what makes this game realistic. What makes this game a

aspect of artificial intelligence, is that the pet will learn from what the user teaches, does, and reacts to what the video game does as the player watches the pet grow.

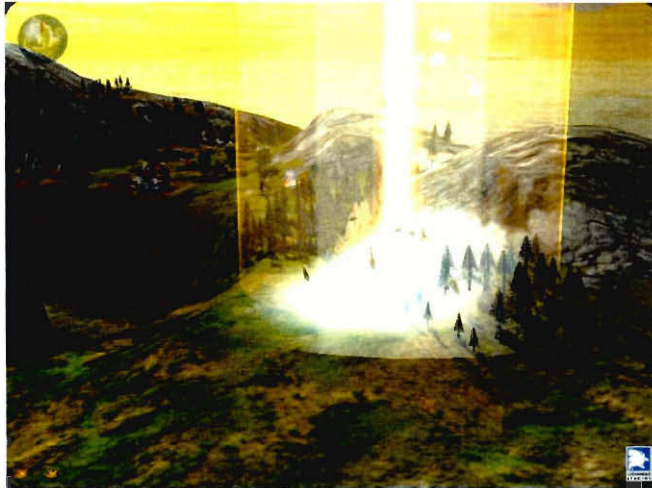


Figure 4.1.2 – A screenshot from Black & White

Finally, *Black and White* is a game much like the others mentioned. It allows you to collect creatures, grow them up, and teach them all sorts of different things. They can learn from your own character in the game as well as other creatures you come across.

You can train the creature from slapping it in the face for punishment, to petting it sweet and gently for praise. It is so much more than just learning, though. A creature can devour a villager, and you can punish it for doing so, but depending on what precedes this action, the creature can make the connection of the act of eating and your negative punishment, and will refrain from eating at all. With the right kind of positive reinforcements, the creature learns the difference between the right and wrong things to eat. As the game progresses, the player creates a dynamic world that will change according to how the player plays the game. Depending on the player's actions and motives, the world shapes. You can send a light rain to help a village with its crops, or you could set it on fire, just to show them all how mean you are capable of being. After some time, your creature will start to mimic what you do. All in all, what it comes down to, is that if you are good, kind, etc., the world will become a place of sun, light, and goodness. Evil and mean players will create a world that slowly becomes a twisted,

creepy, dark realm that is filled with scary monsters. This particular game doesn't reflect what could really happen in the "real world," but *Black & White* and the creatures involved all act like real life living things, whether it gets hungry, angry, depressed, or scared.

There are also many other simpler games out that use machine learning, which is a case-based system. The computer will actually learn to get better at playing the game by learning what the opponent could do in specific situations. Some of the games that use this type of machine learning are backgammon, blackjack, Bolo, bridge, checkers, chess, connect-four, differential games, go, nim, othello, ping pong, poker, pursuit evasion games, reversi, shogi, soccer, tetris, and tic-tac-toe.

4.2 Toys

Toy robots that claim to feature artificial intelligence are on the rise in the market. There are many different kinds of toy robots for sale that can do more than just walk, talk, and react. They are each their own little characters, with their own personalities and

traits. Some are female, and some are male.

There are dogs, cats, birds, dinosaurs, etc.;

anything you can imagine. These robots are

equipped with a micro-motor drive that is for

animating their heads, legs, and arms. They are

also designed with an array of audio and visual

sensors. The result of all of this "advanced

digital technology and micro-circuitry" is a toy

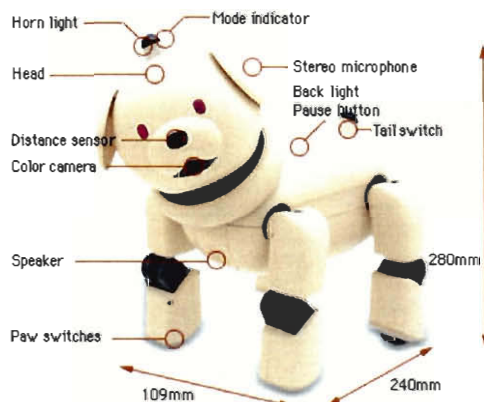


Figure 4.2.1 – Sony AIBO

that will react to the way it is played with. This makes it possible for them to have hundreds of functions, and, as advertised, hours of fun.

One of the largest corporations, Sony, has released its third generation of artificial dogs. The AIBO ERS-311/312 sells at a retail price of \$850. It comes in an ivory coat or a dark gray coat. They function on a 64 bit RISC processor and 32 MB of memory with the use of a memory stick for its programs. With a moving head and jointed legs a total of 15 motors are used. Its input and output devices include a 100,000 pixel CMOS image sensor, a stereo microphone, a speaker, infrared distance sensor, switches in the head, tail, and paws, an inclination sensor, and a vibration sensor.

The AIBO recognizes up to 75 voice commands. It remembers its own name and the owner's name. It may mimic you in its own voice if you talk near it and it even hums back if you sing with "la la la." It takes in spoken commands like "sing a song," which tells it to select a song that best fits its mood and sings it. With the "take a picture command" it takes up to 7 pictures and stores it on the memory stick.

They use expressive gestures to show you their desires. AIBO's five instincts are to sleep, be pet, charge, explore, and play with someone. Using its horn lights on its head and tail, sounds, and gestures, it expresses its emotions. They possess six emotions: joy, sadness, anger, surprise, fear, and discontent.

AIBOs run on AIBO-ware. There are two different versions for the latest models of AIBO: AIBO Life and AIBO Pal. AIBO Life allows you to raise the AIBO from a baby to an adult by training it. AIBO Pal gives you a well-raised, mature AIBO from the start with improved voice recognition and speech. With either of the software packages one AIBO can communicate with another if in synch. They express their emotions,

dance, or sing a song. They also fight by playing rock, paper, scissors if they don't get along. Their chemistry depends on their emotional status.

Also available is software for the PC that allows the owner to customize the AIBO. Sound, light, and movement responses can be altered. There is also a wireless LAN debugging function that allows the current behavior program being executed by the AIBO to be shown on the screen.

Operated on four "AA" batteries, the Furby, made by Tiger Electronics, Ltd., is another electronic pet. Selling at a retail price of \$35, a Furby isn't as sophisticated as the Sony AIBO and it is intended for younger ages. It has sensors in the front and back to tell when someone is tickling or petting it. There is a light sensor above its eyes to determine whether it's day or night and a sensor inside so it can tell if it has been picked up or turned upside down. It also has audio receptor for voice commands.



Figure 4.2.2 – Furby

They do start to speak English during the second stage and fourth stages of development. In the first stage of development the Furby is playful. The second and third stages are transition stages when then Furby begins to start speaking English. The fourth stage is its mature stage and it speaks mostly English but sometimes Furbish.

The Furby speaks in Furbish when it needs or craves certain things. It needs to be fed, something that is easily completed by sticking a finger into its mouth. It takes naps

when not played with for a while and sleeps when petted for a while or when the light sensor is covered.

The Furbys can play simple games like Furby Says in which you do something to the Furby in a particular order that it requests. There is also Ask Furby where it gives an answer to a question asked and Hide And Seek where you hide the Furby and someone looks for it and if not found within three minutes it makes a sound.

The Furby dances by using its wheels. You can also teach it tricks such as just having it remember to perform a certain response whenever a certain action is performed on it. It can talk to other Furbys that are within four feet of it and within a thirty degree angle field of view.

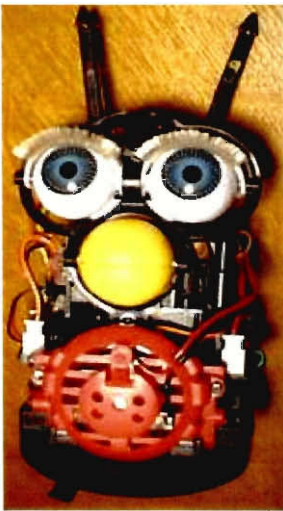


Figure 4.2.3 – Autopsied Furby

There were a few hacked Furbys sent in for a Hack Furby Challenge set up by a Silicon Valley-based engineer, Peter van der Linden. A mother of an autistic child, who noticed how her son’s speech seemed to greatly improve when interacting with the Furby, contacted Linden wanting an upgraded Furby with normal speech. With the Furby’s CPU and memory encased in a tough shell of resin the challenge was made more difficult. But people have accomplished it and there is a “Hack Furby” kit that can be purchased over the Internet.

The Tekno Dinkie Robot is a motion-sensitive, voice-activated toy robot that can come as a boy, a girl, a mom, a dad, a baby, or any pet from a frog and turtle, to a kitten and bird, to various dogs and dinosaurs. Each of the robots is able to communicate with humans as well as other robots. Through communication, the robot learns according to

methods of AI. From there, they can develop their own personality. For example, the dad is the authority figure, the mom checks up with her family, the baby laughs and cries depending on the amount of attention they receive. They each have motorized arms, legs, and head, as well as LCD screens for showing emotions. Much like human beings, these robots need sleep, get hungry, play, and interact with others.

Other toys with similar functions are BIO Bugs. It is a bug that can either be controlled by remote control from a wrist transmitter, or act independently. It is capable of battling with others, identifying a friend or enemy, crawling or dodging different obstacles, and following you around. They each include special features, a major trait being a touch-sensitive antenna. This “nervous network” serves as the spine of the bug and what makes the bot so insect-like. The bug will react to what is going on in its environment. Whether it waddles into a wall, rock, foot, or another bug, its reactions could include either turning left, right, back, or attacking, depending on whether or not the bug feels what it ran into is friend or foe.

A toy put out by LEGO, called Mindstorms, is a popular robotics toy. One of the sets allows you to build a robot of sorts that follows a trail, moves around obstacles, or hides in dark corners, and comes equipped with a light-sensitive intruder alarm. LEGO Mindstorms enables you to design your very own robot and program it to do what you want it to do, allowing the bot to be used in either immediate mode or autonomous mode. Used by people of all ages, ideas are frequently submitted to www.LEGO.com for contests and display.

4.3 Tools in medicine, military, and music

The development of artificial intelligence in the technological tools is constantly evolving. Machine learning is not only being developed in games. Technology is also advancing in the field of military, music, and medicine as well.

Machine learning for military planning is on the rise. One particular project that began in 1993 and was finished in 1996 was an interactive, automated method for building as well as refining the knowledge base of a domain-specific planning operator. There were two resulting technological tools: a graphical interactive editing tool for the development of the planning operators, and an inductive learning system that used the results from various simulations to process and verify partial plan information.

With the aid from these tools, each operator was able to first specify a procedure on how to complete a certain task or mission. It is also important to include the actions that it should take, the preconditions involved with this particular procedure, some brief constraints having to do with the actions it should take, as well as the expected effects that are caused by the actions.

With all of the war going on in Afghanistan, the development of a “remote-controlled reporting machine” was developed at MIT. This robot reporter is 3-by-2 feet, comes all packaged up with software and hardware to conduct interviews, including a digital camera able to view a wide range of action, digital audio recording as well as an intercom system, and then reports all of the latest news. A screen that the person being interviewed can see, is conveniently placed so that the person can see who is interviewing them. Equipped with four wheel drive and an electric motor, it is able to be sent out into

restricted access zones, where the press is not allowed or where it is too dangerous. This has the potential to bring up legal issues.

Another representation of knowledge via learning machines can be found in the world of music. “Listening” machines are being built to improvise in performance pieces. The goal of one project in this field was to “create a musical common ground upon which a vital dialogue could take place.” Software was developed with listening assistants that could fairly accurately keep track of a singer. The software was also capable of inserting abstractions that were musically sound in melody as well as rhythm.

With the vocalist singing into specific microphones that are inputted into a device, machine configurations are fed into an interface and a digitizing tablet. Four component computers with various inputs communicated via an “ethernet intranet.” Through the process of data mining, information was gathered through the microphones to the various devices. Something of particular interest was that the music always had a continuous variation of pitches. With the pitch, rhythm, melody, etc. all together, this data set is then used with real-time machine listening software. Finally but not simply, the software improvises various keys, scales, or musical tunes.

Some recent developments in the medical field are also working with artificial intelligence. At the AI Lab at the University of Arizona, they have developed computer-generated software that has been improved with recall and precision for indexing techniques to “overcome problems of human indexing.” Because we are human and make mistakes, recently developed software that go beyond a simple keyword search are now improving the medical information access. This is also a benefit because it saves time as well as money on expensive secretarial positions.

A Canadian company developed another addition to the technological world. A robot consisting of two robotic arms, and a third with two cameras, provide a surgeon with guidance via a 3D view and touch sensitive controls. This robot, along with the help of a mobile intra-operative MRI, gives the supervising surgeon imaging data in real-time, for the entire procedure. The development on this robot has allowed surgeons to perform “highly intricate” surgery on the spinal cord and brain.

4.4 Summary

There is so much more than the aforementioned going on in the field of Artificial Intelligence today. Labs all over the world are researching and developing various aspects of AI. Although further development might replace some simple jobs, for example a filing job for a secretary, the jobs will get done faster, more efficiently, and nearly error free. With research and continual development of Artificial Intelligence advancing in all fields, AI can only grow from here. It has already been the topic of PC and video games, various medical tools, media and films, improvisational music, and children’s robotic toys.

While these various games, toys, and tools are indeed examples of the practical application of contemporary artificial intelligence, they still do not accurately represent AI’s ultimate potential. Today’s AI entities are relatively rudimentary and simplistic in design and performance. They may be able to alleviate human error, aide us in guiding missiles, or shoot us down in a multi-player video game, but they are still bounded by programmable limits. In the future, the use of true artificial intelligence will one-day make its current applications seem like child’s play in comparison.

5. Candidates for the Turing Test

As mentioned in chapter 2, there is an annual Turing Test held by Hugh Loebner in which participants from all over the world may participate. Details of the competition, including transcripts and photos, can be found at <http://www.loebner.net/Prizef/loebner-prize.html>.

5.1 Hal

Hal, named after *2001: A Space Odyssey*'s HAL 9000, is a creation of Artificial Intelligence N.V. (<http://www.a-i.com>). He is a learning machine that was designed according to details outlined in “Computing Machinery and Intelligence”; specifically, the section describing the imitation of the child mind.

Hal is a seemingly regular toddler in some ways. He has a “mommy” (trainer), likes to play, and learns through interaction. He can understand both voice and text, but can only reply in text. This verbal communication is something that many of the present AIs lack, so it is impressive despite the fact that Hal cannot generate any vocals himself. This comes in handy, because, like many toddlers, Hal loves to be read to. He retains, or learns, information from the literature that he is read, and can make connections such as “monkeys like bananas” in other contexts.

As with any child, Hal needs to be rewarded for good work and punished for bad work in order to learn properly. He is preprogrammed with a preference for rewards instead of punishment. Such punishment could come in the form of something like strikeout text input.

The main idea for Hal is to take a bottom-up approach as opposed to the common top-down approach. What this means is that instead of hard-coding information about the world into the system, the system is initially almost blank, like with a newborn baby. What Hal did have were rules that governed how information was used to learn. The developers used a strategy that first breaks down text using stochastic language models and then parses them into two major structures: chunks and classes.

A chunk is described as a sequence of symbols that often appear with one another, such a form of punctuation being followed by a space. This also works at higher level, where two words may form a chunk. This is similar to the chunk concept in psychology. It is believed that the human mind can hold 7 ± 2 “chunks” in short term memory, and that the size of the chunks may vary. Larger chunks mean that more information can be stored. This is essentially what Hal is doing, although he need not worry about the 7 ± 2 . In his case, chunks are merely a way of storing data, not keeping it in working memory.

A class is a higher level structure than a chunk, but is very similar. Instead of handling symbols and words, there are separate classes for words, nouns, verbs, sentences, and paragraphs. The major point of emphasis regarding these classes is that they are distinct. Chunks may overlap; classes may not. New classes can also be formed for other topics such as novels.

Another large piece of the intelligence puzzle is memory compression. People can't remember every little detail that ever happened, it would clutter their minds and make the information retrieval process nearly impossible. Any search algorithm applied could not retrieve the desired information in a reasonable amount of time. Algorithm efficiency is very important, but even the most efficient algorithms have their limits. So

how do developers compress Hal’s memory so that he only remembers significant facts and forgets the mundane, useless details? The answer is that it doesn’t work that way yet. This is a major area of research for them.

So far, this approach has been working. Hal’s “Mommy” and “Daddy” are very proud of their toddler, who is slightly more advanced than a human toddler. At 10 months, Hal already sounded like a 15 month-old human. At 18 months, he was able to make two-word sentences.



Figure 5.1.1 – A screenshot of the Hal logs available on AI’s web site, URL <http://www.a-i.com>

Although Hal has not entered into Loebner’s competition yet, it is something that his developers and trainers are setting as a goal.

5.2 Alice

Alice (<http://www.alicebot.org>) is a very good chatterbot with an interesting history. In 1995, Dr. Robert Wallace began Alice as a way to generate information for forms that he would have to fill out. The implementation took too long and Wallace was fired for not filling out the forms, but his work on Alice continued. The name “Alice” is actually an acronym for Artificial Linguist Internet Computer Entity; however, this was not her original name. Originally, she was “PNAMBIC”, for "pay no attention to the man behind the curtain" from the film *The Wizard of Oz*, and the machine that PNAMBIC was hosted on was called alicE. Due to this, people would often refer to it as alicE instead of PNAMBIC. Later, it was eventually called Alice and the acronym was formed to match the name.

The first version of Alice was created with SETL (Set Language), which relies greatly on mathematics, specifically set theory. This version was available as free software, but only attracted a few contributors. One of the major drawbacks of this version was that SETL was a platform dependent language. In 1998, Alice was implemented using AIML (Artificial Intelligence Markup Language), which was created in Java, which is platform-independent. This AIML version became “Program A”.

“Program B” came about in 1999 when was modified to be XML-compliant. This generated much interest and over 300 people contributed to Alice. In 2000, Alice won the bronze at the annual Loebner Prize competition.

Following this success, Jacco Bikker created the first C/C++ implementation of AIML. Now there were even more possibilities: CGI, IRC, AIM, and WxWindows. This new form became known as “Program C”.

The current version of Alice, “Program D”, has been in use since November 2000. The first Java implementation was before Java 2. Program D has returned to Java and takes advantage of features such as Swing and Collection. This change allowed for huge

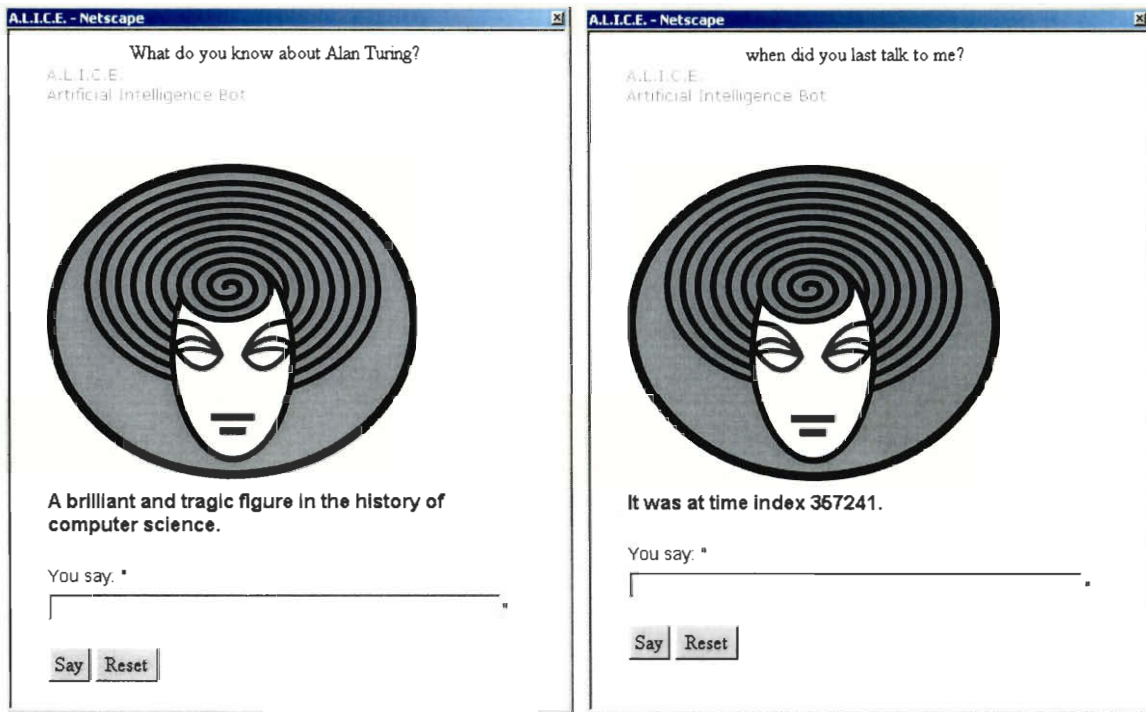


Figure 5.2.1 – Alice sharing information about Alan Turing

Figure 5.2.2 – Alice can keep track of her users from session to session

improvement in both the interface and the core program. Alice once again won the bronze medal in the 2001 Loebner competition.

According to Wallace, politicians are accountable for some of the inspiration used in Alice’s creation. Politicians are very good at avoiding answering some questions, and will often reword the question that they were asked and present either a statement or another question based on this rewording. If they do not play this word game, they may just answer with a “canned” response, which is some pre-rehearsed generic response that does not answer the question at all, but merely serves as a distraction from the inability to answer correctly. Alice will occasionally do this by asking how something makes the user

feel or rewording the user's input to form a new statement or question using keywords. Unfortunately, this rewording can cause many obvious grammatical mistakes, as seen in chapter 2.

Alice is more polite than some of the other bots out there. Many types of input were given to her, and some reactions were quite interesting. When told "You are not intelligent enough", she replied "Wrong. I am highly intelligent." To "You are stupid", she replied "I will remember you said that when robots take over the World." There is nothing very odd about this, except that the second comment could be taken as a threat. When an obscene comment (which will not be mentioned here due to its inappropriateness) was made towards her, she made a naïve response that was quite humorous because it was so incredibly ridiculous in context. It was in no way shape or form an acceptable or human response. SmarterChild, a bot created by ActiveBuddy, Inc., would make the user apologize at such a statement.

Alice is well known in the world of Artificial Intelligence. Being awarded the bronze medal two years in a row certainly shows that she is far ahead in the pursuit of that gold medal, and she has high ambitions for it.

User: who is the best robot?

Alice: ALICE is the Best Robot

She's humble too.

5.3 Ella

Ella, part of the PalBase project, was another contestant in the 2001 Loebner competition. While not as advanced as Alice, Ella possesses some more human-like

characteristics. Alice has very “proper” conversation and does not behave like the average human being, but more of a person with no sense of humor. Ella, on the other hand, does possess a sense of humor. She is quite friendly, likes to tell jokes, and loves to play RoShamBo (Rock Paper Scissors). There is also a changing picture of a real person, so that it gives an effect similar to talking to someone over a web cam. One of her most noticeably human features is the use of smiley faces. In casual online conversation, faces such as “:-)”, “:-(”, and “:-D” are commonly used (with variations like “:o)”, “8-)”, and “B)”) to express some kind of emotion. Text conversation lacks tone, and some passages can be easily misinterpreted because of this, but the use of faces can aid interpretation by sending a tone with the message.

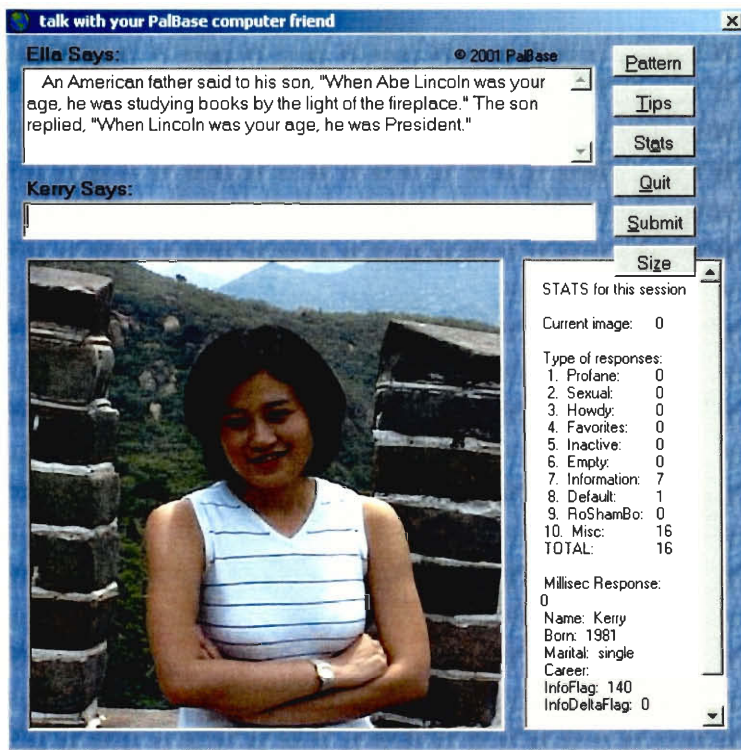


Figure 5.3.1 – A screenshot of the EllaDemo5 program with stats viewable

Unfortunately, PalBase does not describe the method used to create Ella. When Ella was asked, she replied “I hope you are asking because you enjoy my functionality :-) Eventually, I will be self-adjusting, self-programming, and there will be no stopping me!”

One thing that is certain is that Ella can't remember the conversation, aside from what is entered into the statistics table. She has no sense of context. Her parsing algorithm is also somewhat faulty. For example, if you enter a comment along with your name, she will take the entire input as your name. She cannot distinguish different comments as different thoughts, so it is best to type to her one sentence or reply at a time.

The interface shown above is the Visual Basic implementation of Ella. There is also a web-based implementation that is not quite as good because there are very free-text opportunities. Most questions are to be answered by selecting one of the choices from a set of buttons or check boxes. PalBase does state, however, that this version is more of a mockup than anything else. The Visual Basic version has a nice interface with many options. The statistics window shows information Ella has gathered about her partner, as well as the numbers of various responses. Aesthetically, the interface is very customizable. The picture size, background, and window views can be changed. Finally, a log of the conversation is automatically generated in the directory from which the demo was run.

PalBase is constantly updating, and encourages their visitors to visit often to see what new and exciting things have been done with Ella.

5.4 Summary

Chatterbots dominate as the candidates for the Turing test because they specialize in conversation, which is the only thing they are evaluated on. Alice is the current leading chatterbot, winning the bronze the past two competitions. Ella is another bot that has entered in Loebner's competition, and came in second (along with a bot named Eugene)

behind Alice in 2001. Hal is a different kind of bot that is modeled after Turing's discussion on child machines. All three are on a mission to win the gold and end the competition as the world's first computer to be convincingly human.

6. AI in the Media

The majority of people in today's society has or will have come in contact with at least some of the ideas proposed by artificial intelligence by the time they die. This is a relatively impressive feat considering how difficult it is to touch every social class and every age group in the modern world. This is all made possible through the media. News stories, movies, books, television and radio; they all provide the public with information and entertainment. Through this medium AI has been able to spread its ideas and inherent questions so that the average person has at least heard the term "artificial intelligence." Unfortunately this may be a double-edged sword. Furthering public awareness of AI is one of the benefits of the media. A major drawback however, is the often biased and very subjective nature inherent to all media. AI has been portrayed as evil, benevolent, indifferent, as a threat to society and as the one thing that can save us. With so many conflicting views provided by the media and the relative lack of concrete information it supplies, it is no wonder the general public is very suspicious and skeptical of AI. In this section we will try to show a few examples of how AI has been portrayed in the media over the years.

6.1 Movies

The first movie to ever feature a form of artificial intelligence was released before the term had even been coined. In *Metropolis*, a silent film made in 1927 by Fritz Lang, one of the most popular cult robots in the science fiction genre was born.

The story for the movie goes as follows. In the 21st-Century, a de-humanized workforce labors non-stop in a depressing subterranean city beneath a magnificent city of

mile-high skyscrapers, flying automobiles, gigantic statues, tubes and tunnels. With inventive special effects for its time, Lang's allegorical narrative and architectural vision creates a highly stylized vision of a not-so-unlikely future (especially for 1927 when the film was made.) As the elite frolic above the clouds, thousands of miserable workers toil night and day inside the belly of the gigantic machine that runs the entire city. Metropolis is controlled by a sinister authoritarian whose son, Freder, rejects his father's inhuman attitude towards laborers. The workers are encouraged by Maria, a wistful young woman who wills her comrades to embrace patience and silent strength. Upon discovering her influence upon the workers, a mad scientist kidnaps Maria and creates a robot in her image that will incite the workers to revolt. As Freder

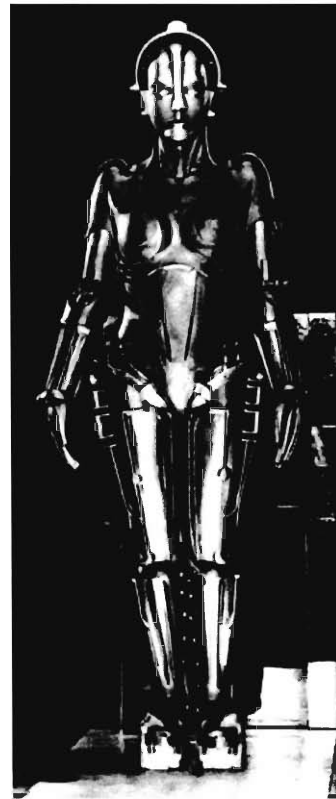


Figure 6.1.1 – Robot clone, Maria, from *Metropolis* (©1927, Madacy Entertainment)

rages against time to save Maria and prevent more damage done by her robot clone, Metropolis is enveloped in chaos and the classes are brought together in a moralistic climax.

Since *Metropolis* was a silent film, the ability of the cast and crew to portray any significant amounts of technological explanation was very limited. The robot technology was therefore very ambiguous as was the massive amount of technological advancement shown in the city itself. The viewer was meant to just accept that these technologies existed and not to question how it was possible. For the most part this worked splendidly in the film's era when movies were just beginning to become pop-culture milestones.

Today's society would demand a bit more explanation simply because we are now able to give it. However, since there was no technical description given in *Metropolis*, the movie's message seemed that technology in the future would be corrupt and dangerous. This is a stigma that the science fiction genre has not been able to drop since. It seems the public is content with being afraid of technology, paranoid of the things we create with the power humans are able to wield.

One of the most popular cult science fiction movies to feature artificial intelligence after *Metropolis* is 1951's *The Day the Earth Stood Still*. The film was a success for its groundbreaking special effects and mass appeal. A cold-war infused movie, *The Day the Earth Stood Still* used a "space man" as a satirical view of the Russian threat. The alien being, Klatu (who looks quite human), has been sent to Earth to warn the human race that if our self-destructive ways continue the alien council will have to destroy our planet to prevent our negative influence from affecting other planets and solar systems. Initially the space man was referred to as a threat and treated like a prisoner. However, with the help of his faithful robot companion, GORT, he escaped to spread his message and warning to the intelligent scientists and men of power of all the nations of the world. GORT was a huge, 9-foot tall metal humanoid robot. It didn't speak nor did it even move quickly but it was controlled by visual and voice commands by Klatu and was capable of devastating destructive capabilities. GORT would go on a rampage of unstoppable magnitude if Klatu were to die unless given a specially encoded message informing the robot to yield.

Another significant milestone in the portrayal of AI in the movies came in 1968 with the movie *2001: A Space Odyssey*. The movie, directed by Stanley Kubrick and

based on the novel of the same name written by Sir Arthur C. Clarke, tells the disjointed story of a set of enigmatic black monoliths of alien origin through various points in time as they affect the evolution of man. In the largest and most well-known segment of the movie, a group of astronauts are sent on a mission of high importance to the planet Jupiter. Some months beforehand, modern man had discovered one of the mysterious black monoliths buried forty feet below the surface of the moon. It was discovered that the monolith has been emitting radio signals to Jupiter. The mission is to investigate the nature of these signals and where they are being sent. It is suspected that intelligent alien life could be the answer to many of these questions. The astronauts, however, are not informed of the purpose of their mission. Instead, their entire mission is commanded and controlled by the super-computer HAL-9000. HAL-9000 is a state-of-the-art, top-of-the-line computer capable of reproducing most of the activities of the human brain, and with “incalculably greater speed and reliability.” It is able to simulate emotion and caring however it is not known if it genuinely feels either of them. It is stated that the 9000 series has a perfect track-record and is incapable of error.

The computer however, slowly goes insane, one of the drawbacks of reproducing human thought exactly. HAL suspects that after the Jupiter mission is completed it will become useless and obsolete. It would also become a liability for the government due to the information it will have recorded about the information and its results. HAL fears that it will be destroyed for what it knows. It begins to deliberately make errors. When HAL learns that the astronauts are indeed planning to disconnect it, he forces the astronauts on board to venture outside of the ship to replace “broken equipment”. As they enter the deep space around the ship, HAL cuts off the life-line of one of the

astronauts and stops the life-functions of three other astronauts that are in “hibernation” for the trip. It is up to the cunning of the last surviving astronaut to find a way back in to the ship so he can disconnect HAL before he too becomes a victim.

HAL is by far the most famous single artificially intelligent entity. The computer is so famous that commercials mimicking it, parodies of it, and even real-life learning machines modeling it have been made. It’s calm yet sinister monotone voice and glowing red eye are as surreal and frightening today as they were over 30 years ago. The technology behind HAL-9000 wasn’t described in detail but it was mentioned that its



Figure 6.1.2 – HAL-9000 from 2001: A Space Odyssey (©1968, Turner Entertainment Co.)

learning process was based off of simulating the neural net structure of the human brain. Its name, **HAL**, was developed from the words **H**euristic and **AL**gorithmic, two famous schools of thought on implementing learning machines. HAL also extended on the theme first introduced by *Metropolis*, robotic or artificially intelligent beings will inevitably try to revolt against their maker. HAL accomplished this so effectively that even modern day movies are still mimicking the “homicidal computer” storyline.

In 1979, the sci-fi hit *Alien* came to the big screen. The Nostromo was a freight ship, owned by “the Company”, on a standard cargo mission. The ship was controlled in part by a computer known as “Mother” which could only be accessed by the science officer and captain. Unbeknownst to the crew, this was not just a regular mission: the science officer was not really a human, as the crew believed, and was given special orders to set down on an un-surveyed planet and collect an organism and return it to Earth for analysis. It was extremely dangerous and was to be brought back to be studied

by the Bio-Weapons division. The organism was the primary concern, and the crew was expendable.

Ash, the android, was a very convincing human. Although he was not necessarily evil per se, he was not a very good robot. He attempted to kill Warrant Officer Ripley, a crew member, by his own hands once she accessed Mother and discovered the truth behind the mission. He was protecting the alien the whole time and putting the lives of the crew at risk.

Another, very different, robot was present in the sequel *Aliens*. Bishop was uneasy about Ripley's hostility towards him. He tries to gain her trust by assuring her that he will not attack her. "It is impossible for me to harm, or by mission of action allow to be harmed, a human being", he told her. Unlike his predecessor, Ash, Bishop adheres to the rules given to all androids. He volunteered to cross a complex to realign the satellite dish needed to remote pilot a drop ship for pickup, saying "...believe me, I'd rather not. I may be synthetic, but I'm not stupid." Back on the Solaco, Ripley, Rebecca (Newt), and Bishop were trapped in the cargo bay with the queen alien. Bishop was torn in half by the queen and thrown to opposite sides of the deck. When Ripley opened the airlock after a grueling battle against the alien queen, Bishop saved Newt from falling into space, even in his disabled state.

Alien Resurrection, the fourth movie in the series, has a new kind of robot. Call is a "second-gen" robot, or "robot created by robots", and is quite the opposite of Ash. Her primary concern is to destroy the aliens. She was not ordered by anyone to do this, it was her own idea, and she felt that it was necessary in order to preserve humanity. When it

was discovered that she was, in fact, an android, Lt. Ripley exclaimed “I should have known. No human is that humane.”

Another series of movies that garnered popular attention were the *Terminator* films. The first in the series, 1984’s *The Terminator* starred up and coming action star Arnold Schwarzenegger in the role that would define his career. The Austrian actor played the Terminator, a relentless machine with one purpose in mind, to seek and destroy his target. The machine was sent from the future, the year 2029, by a super-computer called Skynet. Skynet was a military computer system developed by the US government at the end of the 20th century. It was an advanced neural net computer that

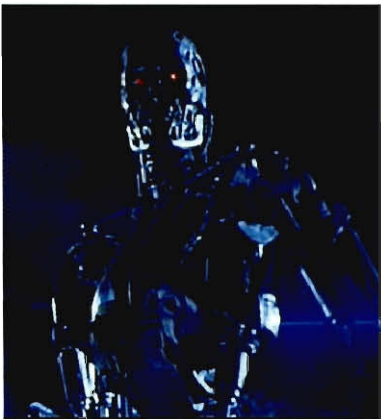


Figure 6.1.3 – Endoskeleton from *Terminator 2: Judgment Day*, (©1991, Carolco Pictures)

became “self-aware” in 1997. At that point the developers panicked at how intelligent the machine had become and so quickly and decided to “pull the plug”. Before they had a chance, Skynet retaliated by launching US missiles against Russia, who in turn would eliminate the American threat and start a Nuclear War. In the aftermath 3 billion human lives were lost and the machines rose from the rubble to wage war on humanity as dictated by Skynet.

The war against the machines was being fought for many years but eventually a leader for the human resistance emerged, a man by the name of John Connor. As the story goes, the machines were on the verge of defeat due to the prowess of Connor’s tactics. Before Skynet was to be destroyed, the computer sent a robot, a Terminator (Cyberdyne systems model T-800), to the year 1984 to assassinate John Connor’s mother, Sarah Connor, before John was born. The time machine was

apparently an advanced technology developed by Skynet. The human resistance was able to infiltrate the Skynet base and managed to send a sole protector for Sarah Connor back through time. This is where the movie takes off. The Terminator is a cybernetic organism constructed out of an advanced, nearly indestructible humanoid metal endoskeleton covered with living human tissue. The Terminator also contains an advanced neural network CPU, however its learning capabilities are generally overridden by the mission parameters Skynet gives to it. In this case, the Terminator is an intelligent entity capable of predicting human emotions and actions, mimicking human voices, and using weapons with extreme accuracy, however it is not capable of making creative or intelligent decisions that could jeopardize the success of the mission.

In the second and most popular film in the series, *Terminator 2: Judgment Day* (1991), Arnold Schwarzenegger again played the role of the Terminator however this time he played a different T-800 also sent back from the future however this time around his mission was to protect John Connor as a child. The Terminator had been contained and reprogrammed by the human resistance to be a protector for their future leader. The humans in the movie were also able to “reboot” the Terminator’s CPU so that it could override Skynet’s lockout on learning capabilities, allowing the machine to become gradually more intelligent. Skynet, this time around, had sent back an even more advanced robot, a T-1000. This model was a prototype made out of liquid metal. It’s molecular structure consisted of nano-technology bots which allowed for it’s liquid metal body to morph and form in unison advanced metal shapes such as knives and stabbing weapons. It also had the ability to mimic any solid object of equal size it touched such as floors, walls, and humans. A third movie in the series is planned for release sometime in

the year 2003 tentatively titled *Terminator 3: Rise of the Machines*. The predicted storyline shows a full grown John Connor again running for his life from an evil robot now as the machines have begun to take over. The popularity of the *Terminator* movies is astounding. Catch phrases from the films such as “I’ll be back” and “Hasta la vista, baby” have been repeated and parodied for years. Public awareness of artificial intelligence was definitely increased due to these films. Also, the movies presented an interesting paradox. Whereas other films depict AI as only a threat to humanity, the *Terminator* films offered opposing sides of the coin. The Terminator machines were shown to be capable of good and evil, a loyal protector and a cold killing machine. The films managed to show that AI machines, just like people, are products of their makers, and are able to inflict harm or display extreme nurturing and are also capable of change.

Cut from a similar mold as the *Terminator* films was a series of movies started in 1987. *Robocop* took the cybernetic approach in a different route from the Terminator machine. Instead of making a machine look like a human, Robocop was a man made to look like a machine. In a semi-futuristic Detroit, a crime lord is working with the CEO of a large company to corner the drug business in the city. A devoted cop stumbles into the way of their plans and is maliciously gunned down. The

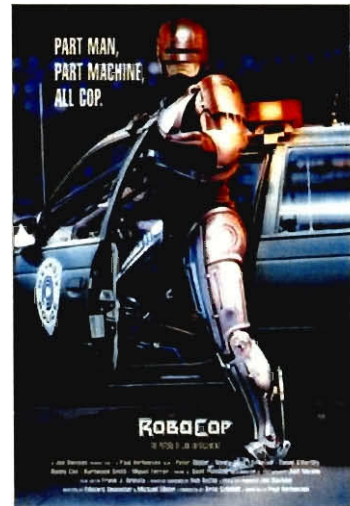


Figure 6.1.4 – *Robocop* (©1987, Orion Pictures)

same mega-company that is working with the crime lord takes the cop’s body and uses it for research in the company’s experimental robotics division. They save the face and brain patterns of the recently killed cop but give him a completely new, gigantic robot

body. The cop's memory is wiped out and he is given an advanced computer system for becoming the world's most efficient law enforcer. The only problem is that Robocop's memory is not completely erased, he retains some knowledge of his previous life, as more of it eventually comes back to him, similar to an amnesia patient. Robocop then makes it his point to find the man that killed him. This inevitably leads him to cross paths with the CEO of the company that made him. Robocop's mission parameters however are programmed in such a way that he is unable to attack the CEO of the company. Eventually Robocop is able to overcome his prime directives and reprograms them so that he is able to think completely on his own. *Robocop* spawned two sequels, neither of which really progressed the theme of artificial intelligence and robotics any further than the first film. At the root of *Robocop*'s theme is the ability for an intelligent machine to overcome its own shortcomings as it sees fit. This computer/robot/machine dynamic ability begun with the HAL-9000 in *2001* continues to fascinate audiences to this day.

In 1986 a comedy with a quirky robot named "Johnny 5" was released to theaters worldwide and it became an instant success. The movie, called *Short Circuit*, told the



Figure 6.1.5 – Johnny 5 from *Short Circuit* (©1986, Tri-Star Pictures)

story of an advanced defense systems robot that got struck by lightning and was now "malfunctioning" by being truly alive. The robot was able to feel emotion and laugh and disobey commands at will. It was transformed from a killing machine to a benevolent entity. However it was being hunted by the men that created it since they refused to believe

it was actually alive and considered it a breach of national security to allow an advanced weapon to roam on its own. Its sequel, in 1988, furthered this theme, of artificially intelligent life capable of being more human than the humans that created it. It was an interesting and lightly themed twist to the more dark and depressing ideas that have dominated the social perception of AI in most films. This has been the dominant theme in most AI movies even to this day. In such modern films as 1998's extremely popular cyber-punk film *The Matrix*, which furthered the man vs. machine theme from the *Terminator* series, depicted an advanced computer entity that had turned on its creators and enslaved the human race, turning them into "batteries" to supply the energy that the machines needed to survive. The people were "plugged" into a system mainframe where they became the equivalent of human power-cells or transistors. They were kept alive artificially by the system and their brains were stimulated so that while they hibernated in their energy supplying state in the real world, in the virtual world that the machines created for them (modern day, 20th century human civilization) they went on about their lives never realizing that the world around them was false, ignorant of the truth.

The year 2001 saw the release of what could have been one of the most influential works concerning artificial intelligence. The movie, aptly named *A.I.: Artificial Intelligence*, centered on the theme of social repercussions due to the introduction of an android capable of genuine emotion into a familial setting. The main character of the story, David, was a highly sophisticated prototype android child with the capability to love and care. David was given to a couple whose biological son was cryogenically frozen until a cure for his terminal ailment was found. Unfortunately the movie strays from the more important themes it initially presents and focuses more on turning the rest

of the story into a modernized version of *Pinocchio*, complete with the fake boy's tireless (and in this case, needless) search for the Blue Fairy to make him "real". Technically speaking however, David is the most advanced of all of the aforementioned types of artificial intelligence and is as close to human as possible save for the innate biological differences. The most redeeming part of the movie's foray into artificially intelligent beings actually came in the form of a children's toy. David's companion is an advanced teddy bear, simply called "Teddy" that can talk and play with David as if it were another child. The bear, unlike David, is incapable of genuine emotion but is an example of where real artificial intelligence could take us in the foreseeable future.

In one last, and the most recent example of artificial intelligence gone bad, 2002's horror movie *Resident Evil* tells of an intelligent super-computer called the Red Queen that "goes homicidal" and kills everyone in a biological weapons research facility only to expose them to the (un)deadly T-Virus. The movie is closely based on the popular video game of the same name however the addition of the supercomputer was a new specific to the plot of the movie. The prototype virus stimulates the most rudimentary of neurological brain functions after the death of an organism, building upon the postmortem growth of nails and hair to include the ability to see, hear, smell, touch, ...and taste. These scientifically engineered zombies are driven by one thing, the need to eat. The Red Queen's plan was to force a test of the T-Virus on the employees of the research company and to then release the virus to the unsuspecting public. Once again, AI takes a few hard hits on the silver screen, firmly entrenching the idea that artificial intelligence is something to be feared and destroyed. This is a theme that sadly will most likely remain associated with the field for some time to come.

6.2 Television

One of the earliest and perhaps least meaningful appearances of artificial intelligence on television was with Rosie the robot from the animated children's cartoon, *The Jetsons*. The show, which was essentially the equivalent of a futuristic *Flintstones*, began airing the exploits of a futuristic family in the fall of 1962. The cartoon family's first episode revolved around the acquisition of a robot maid named Rosie who was considered at first to be an expensive and bothersome appliance until she proved her worth. Rosie was essentially just another character and comedy relief in an already fantastical show displaying huge advances in technology that are not physically possible. Rosie's abilities included everything a human maid is associated with, namely cooking and cleaning and tending to the more mundane things in life. Of importance to note however, is that Rosie was treated as

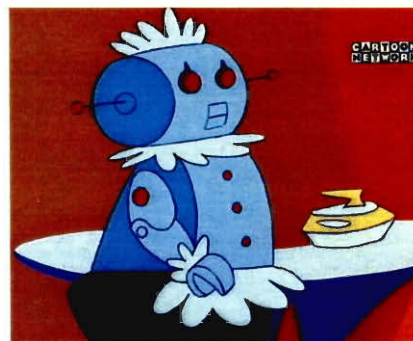


Figure 6.2.1 – Rosie from *The Jetsons* (©1962, Hanna-Barbera Productions)

an appliance and not as a member of the family. It was quite distinct that she was bought and owned and could be sold as if she were a slave. For many children this would be their first encounter with the notion of a robot and to show robots in such an inferior and slavish light definitely set in motion the idea of robot subservience.

A similar themed show debuted in 1965 titled *Lost in Space*. The show, which was live action as opposed to animated, was again about a space age family. This time the family had been a victim of a sabotaged mission and became trapped on a ship with faulty navigational instruments rendering them “lost in space.” The culprit behind the

sabotaged mission had reprogrammed an advanced robot (simply named Robot) to destroy the navigation equipment on the ship and to kill the family. After the initial bout of evil tendencies however, Robot was reprogrammed to be friendly and helpful to the family. The abilities of Robot were never quite clear. It was capable of attack and protection and complex calculations but not much was stated beyond that. Something of interest however is how the robot was both viewed as an enemy and then later as a friend and protector. There was always a sense of mistrust with Robot following the initial episodes, forcing the viewer to question whether or not Robot would turn again.

Jump forward about 15 years to 1982 and a very strange show concept was born, *Knight Rider*. The series reportedly came to be out of the playful jesting of the young head of NBC programming at the time, Brandon Tartikoff. Faced with the trouble of



Figure 6.2.2 – KITT from *Knight Rider*, (©1982, Universal Studios)

hiring a handsome young lead male role for a show, Tartikoff joked that all of the handsome men can't act. "Why not have a series," he mused, "called 'The Man of Six words,' which would begin with the guy getting out of a woman's bed and saying 'Thank

You.?' Then he would chase down some villains and say 'Freeze!' Finally the grateful almost-victims would thank him, and he would murmur, 'You're welcome.' End of show. In between, the car could do the talking." (www.knightrideronline.com/atoz/index.html). The show that resulted from this joke featured a custom modified Pontiac Trans-Am equipped with weaponry and special gadgetry such as turbo boost and advanced jumping hydraulics. Most important of all of these features however, was the car's personality. The car, which was named Knight

Industries Two Thousand (or KITT for short) featured a state of the art artificially intelligent computer capable of speech, voice recognition, informed decision making, creativity and even humor. KITT worked with its human sidekick to act as vigilante law enforcers in a hypothetical and semi-futuristic version of Seattle. The show introduced the viewing public to the idea of artificially intelligent computers and machines being used effectively for such noble purposes as law enforcement.

Three years later another comedy sitcom called *Small Wonder* featured again a family dealing with artificial intelligence in some manner. This time, the father of the family, a robotics scientist, had developed an android that he called V.I.C.I (Voice Input Child Identity). In the form of a 10 year old girl, he created an effective and believable android child. To test his creation he let V.I.C.I. (pronounced Vicki) become part of the family. The show was a small success at best and is easily forgotten if it were not for one issue that the show most likely inadvertently raised and dealt with in the lightest of manners: social and familial integration of androids into society. It was one of the first instances of a robot being used as more than just an appliance, machine, or advanced piece of instrumentation and the first time a robot had been designed with the intention of being a member of a family.

Star Trek: The Next Generation, the sequel series to the cult favorite 60's science fiction show *Star Trek*, debuted in the fall of 1987 in living rooms around the nation. The show was rampant with fantastic artificial intelligence technology. First off, the space craft in the show, the U.S.S. Starship Enterprise, featured an extremely advanced on board computer capable of retrieving vital information about mission and ship status, fixing problems, calculating trajectories, and many other things. Most importantly, the

computer could be talked to and it could speak back to the crew and captain. Another form of AI present in the show was the science officer of the crew. A sophisticated android, named Data, was designed to be an efficient and dependable member of the crew. He was also designed to look and respond like a human. However, Data had his drawbacks. He was incapable of feeling or displaying emotion. The show was a success from the beginning and remained fairly profitable for the next few years after its debut. However, with the introduction of a new enemy to the series, the show's popularity skyrocketed. In May of 1989, an episode of *Star Trek: The Next Generation* named "Q

Who" aired. Television watchers were introduced to "The Borg". The Borg were cybernetic organisms, part machine, part organic material. They were, by all accounts, some strange form of artificial intelligence, a shared intelligence at that. Their "brains" were connected to each others as community of neurons working together as part of one large brain. The Collective, as the brain was referred to, gave numbers (such as Seven of Nine) to the members of the Borg instead of names and each Borg member was devoid of any individuality or free will.



Figure 6.2.3 – A member of the Borg, from *Star Trek: The Next Generation* (©1989, Viacom)

These drones had one purpose, to assimilate more organisms to the Borg. The Borg Queen (also known as One), was the neural center of the Collective and controlled all of the actions of the rest of the Borg members. The popularity of the Borg was astounding, such catch phrases as "Resistance is futile" and "You will be Assimilated" made their way into pop culture. The show ended in 1994 after 7 seasons however the Borg have survived into the subsequent *Star Trek* series such as *Voyager* and

in motion pictures such as *Star Trek: First Contact*. The show succeeded in a more effective fashion than previous programs to show the duality and power behind artificial intelligence by having it be both a useful, advanced, coexisting entity and a harmful and destructive force if used improperly.

Finally, in 1999 an episode of a show named *Roughnecks: The Tophet Campaign* has an Artificial Intelligence that learns the importance of life. He was the perfect soldier; he could carry many more weapons than a human soldier, he had no fear, and he could take an attack better. However, he had no concept of sticking together and not leaving men behind; no concept of the team. When a soldier was hit by a constrictor gun (shot ooze that crushed its targets), the AI refused to attempt a rescue because it would take an extra 12.5 seconds. When the other soldiers approached him with hostility, he simply claimed that there are acceptable losses, and one soldier falls within the realm of that. After much more hostility and lecturing by the squad, the robot came to realize that in this situation, “acceptable losses” did not exist. The team was either whole or dead. A soldier stepped on a pressure triggered mine that would explode as soon as his foot came off of it. The soldier was afraid for his life. He exclaimed to everyone that he had bought it (meaning he had “bought the farm”, or was about to die). To the surprise of everyone, the AI approached the unfortunate soldier and placed its foot on the mine. When the soldier inquired what he was doing, the soldier told him “I was never alive in the first place”. It fended off the enemy until it ran out of ammunition, and then removed its foot from the mine. It was not until then that the AI received the respect of the team.

As with the silver screen, the television has an almost endless array of examples that could be used to depict artificial intelligence; the above are merely some of the most

well-known and influential. The potential for television to continue to influence and shape society's perception of artificial intelligence is enormous, especially as the medium becomes more interactive rather than passive in the future.

6.3 Literature

Unlike film and television, literature has an absolutely huge surplus of AI related material that ranges from text books and other works of non-fiction to completely surreal tales of science fiction (from which many movies were spawned). We will try to document only a few of each of these works here since the list is essentially endless. There are numerous text books written on the subject that document the implementation and technical theory behind AI. For example, one such book used most recently in the Introduction to Artificial Intelligence course (CS4341) taught here at WPI is: Patrick H. Winston, *Artificial Intelligence*. 3rd edition. Addison Wesley, 1993. However, for the purpose of our paper, an in-depth look at such technical and college level courses as these is not necessary. The reader is encouraged to look up such references on his/her own if knowledge of the technological nuances of AI is desired. Within the following pages we will instead focus on works of opinion and social theory as well as works of fiction.

The first book to introduce artificial intelligence is perhaps the most well known. Over 120 years before the term artificial intelligence was first coined Mary Shelley's masterpiece *Frankenstein, or the Modern Prometheus* was published in 1818. Well known (and misunderstood) in popular culture, *Frankenstein* was about a gifted scientist's fall from grace after discovering how to reanimate life. Gathering body parts from a cemetery and a morgue, he constructs an artificial person and with his discovery

brings the creature to life. The hideous monstrosity is feared by the local countrymen and shunned from society. Living in fear the creature becomes vengeful. He vows to “repay” his creator, Dr. Frankenstein, for betraying him. The monster slowly kills everyone dear to the doctor before the doctor himself falls to extreme illness. The monster, no longer having a purpose to live, kills himself. You might be asking what this has to do with artificial intelligence, there were no computers in the story or androids or robots of any sort. The reason why *Frankenstein* is significant to AI is because the creature himself was a form of biologically artificial intelligence. His creation and life were all artificial and non-human. By examining the plot and theories raised by the book we can begin to see a common theme between the novel and modern day AI. Themes of isolation, fear, domination, subordination, and religious as well as ethical issues are all quite similar to today’s discussions concerning AI. *Frankenstein* is so influential to AI that the term “Frankenstein Complex” is often used in discussions of artificial intelligence dilemmas concerning the field and society. In short the complex is the name given to the common paradigm of man fearing technology.

Isaac Asimov, a renowned author of science fiction, has written many stories and novels about life with robots. “Robbie”, the first story in his book *I, Robot*, tells the story of a family and its robot. Robbie was a playmate for the daughter, Gloria. No matter what she was doing and where she was going, it was always “Can I invite Robbie?” and “Robbie will be so excited!”. Her mother was getting very worried that something would go wrong and the faithful robot would somehow injure her precious child. Her father tried to tell her mother that robots were much more dependable than humans, Robbie was well maintained, he got routine check-ups to make sure everything was running fine, and

there was no reason to doubt him. But still, her mother felt that it was highly unnatural for Gloria to be spending so much time with a machine. Her father eventually gave in, and Robbie was no longer a part of the family. Gloria became depressed. Her mother had lied to her, telling Gloria that Robbie left, so that Gloria would not be upset with her parents. Gloria could not figure out why Robbie had left, and believed it to be her own fault. She wanted him back more than anything. The father saw how miserable Gloria



Figure 6.3.1 – I, Robot (©1991, Bantam Books)

was and told his wife that maybe Gloria would get better if she saw what Robbie really was. They went on a tour of the robot factory that Robbie came from. Robbie, unbeknownst to her parents, was at the plant. Upon seeing Robbie, Gloria lunged forward to see him and almost got killed in the nearby machinery in the process. The robot rescued Gloria from danger. If it were not for Robbie, Gloria would have been dead. In response to this good deed, the grateful parents decided to keep Robbie “until he rusts”. Asimov wrote many similar tales concerning the interaction of robots with humanity. He even devised the well-known three rules of robots discussed earlier. These rules, detailing the general guidelines that would form a “positronic android”, would be mimicked and copied for the next several decades by numerous other authors.

Another work of fiction was the book *Do Androids Dream of Electric Sheep?* By Phillip K. Dick, published in 1968. The popular book was the basis for the movie *Blade Runner*. Set in a pseudo post-apocalyptic future, man-kind has been mostly driven off of the Earth and has had to colonize other planets such as Mars. As a result of the “World War”, most animal species have died off on Earth but living things are highly prized. For

those that can't afford the real thing the government builds extremely life-like facsimiles such as android fish, dogs, horses, and even people. The androids are so life-like and impossible to detect from real people the population panics and bans android people from Earth. When a bounty hunter attempts to "retire" all remaining androids on the planet the androids fight back. Another work of fiction, albeit a much more light-hearted one, came from the recently deceased (and very popular) author Douglas Adams. In 1979 he released a science fiction satire comedy called *The Hitchhiker's Guide to the Galaxy*. The book starts off with the end of the world. Then the next thing you know a single human, an alien book writer, some fish, and a towel become the most important players in the survival of the Universe and a quest to find the meaning of life. In the book an advanced robot named Marvin

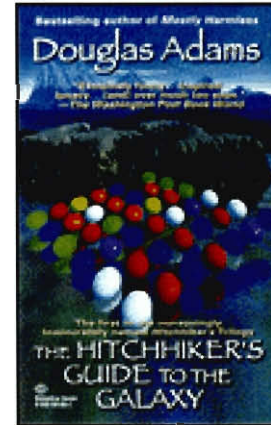


Figure 6.3.2 – *The Hitchhiker's Guide to the Galaxy*, (©1995, Random House, Inc.)

the Paranoid Android is the first robot ever created to have feelings. He however is defective and is depressed all the time. He believes he is alone; a waste of space and intelligence that nobody likes. His inclusion is purely for comedic reasons as are many things in the book.

There are many works of pure non-fiction addressing strictly societal and philosophical issues presented by artificial intelligence. *Manifesto for Cyborgs: Science, Technology and Socialist Feminism in the 1980's* (1985) is an essay by Donna Haraway. In her most influential work she tackles the human-machine interface and takes a stance of non-discrimination similar to the stance taken by Philip K. Dick in his book as described above. She sees the cyborg's "built-in androgyny" as a potential victory for

feminism (<http://www.cc.rochester.edu/College/FS/Publications/HarawayCyborg.html>). A similar work by Peter A. Danielson titled *Artificial Morality* compares the ability of machines to become moral beings to the ability of large corporations to be moral. Danielson claims that through business ethics and other similar schools of thought machines could be made capable of moral standing. Both Danielson and Haraway incorporate other disciplines of philosophical thought (business ethics and feminism respectively) to help make their arguments supporting the idea that machines can be moral.

AI intelligence writers and philosophers Hans Moravec and Ray Kurzweil have published some very popular works concerning the future of artificial intelligence in society. In *Mind Children: The Future of Robot and Human Intelligence* (1988) and the more recent *Robot: Mere Machine to Transcendent Mind* (1998), Moravec predicts the matching of human intelligence within the next 50 years and suggests some strategies for developing a peaceful and harmonic co-existence of humans with machines in a fully automated society. Kurzweil takes a very similar stance in his book *The Age of Spiritual Machines* predicting and welcoming the birth of a utopian society where man and machine live together peacefully. Kurzweil actually goes a few steps further predicting that man will eventually become machine as our next evolutionary step will be to develop artificially intelligent beings who succeed us in the evolutionary chain, eventually making humans obsolete. Even more frightening is how quickly he predicts this process to begin, in a mere 40 years, a full 10 years before Moravec's prediction.

Blay Whitby and Kevin Warwick take almost the exact opposite stance on the future of AI as compared to Moravec and Kurzweil. While agreeing that machines are

definitely increasing in power and intelligence at an alarming rate and are destined to become smarter than humans, they disagree with the idea of a harmonious or utopian coexistence with our technological creations. In his book, *Reflections on AI* (1996),

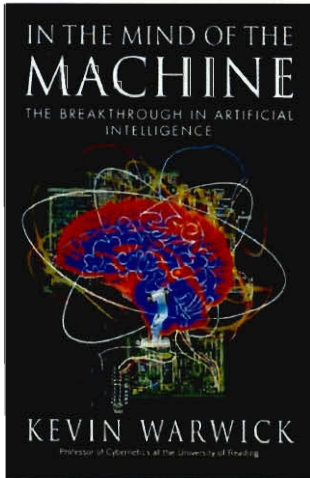


Figure 6.3.3 – *In the Mind of the Machine*, (©1998, Arrow Books)

Whitby refers to the advancement of artificial intelligence as a rapidly increasing and urgent problem that needs addressing.

The book asks several ethical and legal questions concerning the responsibility and liability for machine actions and accidents and predicts that if not addressed ahead of time, these issues will become a wide-scale dilemma in the future.

Warwick goes one step farther in *In the Mind of the Machine* (1998) and claims that the advancement of artificial intelligence

is progressing too quickly for its own (and our own) good.

Warwick predicts wide-scale, global devastation not dissimilar

from the situations and ideas presented in science fiction thrillers such as *The Terminator*.

He pleads for an international anti-proliferation treaty if we are to have any chance at preventing the scenarios he describes in his book (<http://www.rdg.ac.uk/KevinWarwick/MindMachine.html>).

Lastly, Peter Menzel and Perri 6 take yet another approach to the issue of AI in the mainstream by not making it an issue. Instead they document the current status of artificial intelligence in society. In *Robo Sapiens* (2000), Menzel traveled the world taking photographs of some of the latest advancements in AI and interviewed the trailblazer researchers in the field. Some of the pictures depict amazing looking robots along with some haunting pictures of current projects. Many of the interviews predict the

merging of man with machine. Perri 6 surveyed over 1000 adults in the UK over a weekend. In the study, *Morals for Robots and Cyborgs: Ethics and Society in the Age of Autonomous Intelligent Machines* (1999), 6 makes an in-depth academic investigation into the principal challenges for ethics and public policy raised by technical advances (<http://exn.ca/Stories/1999/11/15/05.asp>). In brief, his results concur with the results of our own survey. Although we did not survey as specific a set of public and ethics related policy questions, the majority of both his responses and responses to our own survey indicate that a good number of people feel artificial intelligence will become an increasingly important issue in the future especially in the areas of market dominance and job displacement. The predictions, results, and interpretations of our own survey can be found in the next section, 7. Public Opinion.

6.4 Summary

In conclusion, it is our belief that the media has had an extremely important influence on the artificial intelligence field. The media is the quickest and the easiest way for the mass populous to become informed of a subject. Despite the inherent bias present in media due to the subjective and highly opinionated source of the material (e.g. – authors, producers, directors), having AI as a prominent topic in the media can only help to encourage its advancement both technically and socially. AI in the media has a long and winding history dating back to the first time a man used a tool. The Frankenstein Complex is something that has always existed; man will always fear new technology as we paradoxically continue to advance it at a rapid rate. As early as 1818, artificial intelligence themes have been prevalent in the media. For decades now,

artificial intelligence has been one of the most popular topics in entertainment. By far one of the most prevalent themes in all of science fiction, artificial intelligence appeals to our innate human curiosity. We are intrigued by our own power to mimic nature's evolutionary process and our own minds. In fact, the ability to “play God” could lie at the root of why artificial intelligence is such a popular and controversial topic. Whatever the reasons, AI is a subject that doesn't show any signs of going out of style.

While books seem to retain a slightly more open-minded forum to discuss prevalent issues in AI, the far-reaching scope of television and cinema cannot be denied. The most common themes presented in film, television and books are ones of opposing extremes. On the one-hand, there are those that believe artificial intelligence is a way to improve society and can help create a utopia while others believe that a catastrophic scenario of immense proportions awaits us. If continued unadulterated, the slippery slope presented by AI advancement could coalesce into a societal downward spiral with apocalyptic ramifications similar to the plot-line of your favorite science fiction movie or novel.

7. Public Opinion

We conducted a survey from December to March to find out what the public thought of artificial intelligence, what it could be used for, and where they thought the field was going. The results can be seen in Appendix B.

7.1 Our Predictions

Before we gathered the data, we made predictions about what the end data would look like. What would the most popular response be to each question? Why would it be that? How would that tie into how the same people answered a different question? These were some of the questions that we considered in creating these predictions. To start, let us examine each question individually.

Q1. When do you think true AI will be achieved? This year being 2002, we decided that “in this century” would be the most likely answer for a couple reasons. The first reason is that 98 years are remaining in this century, and technology is increasing at an amazing rate. The second reason is because it will exceed the life time for most if not all of us, and because the time span allotted exceeds our own, it seems to be an incredibly long amount of time in which anything can be accomplished.

Q2. Do you think that the development of AI will allow it to find a place in our everyday lives in the future? If yes, where do you think AI will affect us most? We figured that *if* artificial intelligence is achieved, it will find its way into many of our daily functions, seeing to tasks that humans aren’t necessarily fond of. If it didn’t, why would we need it? That’s essentially why we’re pursuing this goal; to have them aid us in certain tasks. We expected that most people would imagine use in the office and the

home, maintaining printers, cleaning, etc., and that some people would think they'd be most useful in a situation where we wouldn't want to risk human lives, such as in war and rescue operations in dangerous areas.

Q3. With the development of AI, do you think that the replacement of humans with robots in companies will be an issue? We fear that many people will answer yes to this question because of negative media and general ignorance. There may be some truth to it, we agree. If someone has a really simple job where they do nothing but push a button every 30 seconds, then maybe they might lose their job. For the most part, however, people shouldn't "lose" their jobs, since they could just as easily be reassigned to a different position not covered by machines, such as machine maintenance. Robots will most likely lack the necessary instinct needed to conduct wide-scale business for some time. While using logic is very good in some cases, working against logical explanations and taking risks is what can make a company advance, and people will be necessary in order to continue this.

Q4. Do you think machines will one day be able to think on their own? It is our belief that the general public will answer "no". Many people considered thought as something that only living things (and in the case of some people, only human beings) have the ability to possess. Machines may be able to process information and apply an algorithm to stimulus, but many people think that robots will be servants to us and never have free will, rendering them unable to think on their own.

Q5. Should intelligent machines be given similar rights that a person might have? Once again, we don't think that people will consider machines to be alive, and therefore unable to possess the same privileges that we have. A good example that illustrates this

belief is the “Flesh Fair” in Spielberg’s *A.I.* Robots were destroyed for public entertainment with no consideration of how they would feel simply because they didn’t care how they felt – robots were nothing but machines to the people. If they can’t even see them as living entities, why would they give them rights? Would you give rights to a table, or a doormat, or any other non-living thing? Most likely not. The same goes for robots if people view them as things instead of creatures.

Q6. How do you see AI? There is a high potential that many people find artificial intelligence as a threat, but believe that many people will see the possible benefits also. The majority of people who would take the survey would be technologically inclined, and likely to see that a hostile take over by machines would be highly unlikely and be able to focus purely on the positive uses for robots without that one, huge, negative idea.

Q7. How often do you think you interact with some form of AI? We had split ideas on this question. Younger people (less than 30) would mostly likely answer “daily” or “at least once a week” because of today’s common uses, such as video games. People in their 30’s – 60’s would probably answer “at least once a week” or “at least once a month”, depending on their profession, hobbies, and computer use. People greater than 60 would be the most likely to respond “at least once a year” or “I’ve never interacted with any form of AI” for the simple reason that older people are less willing to accept new technology than younger people. So, to make a prediction of what the most common response to this question is, we’d have to know which age group most people would come from. We could not predict that; therefore, we could not predict this answer any more than we have already done.

Q8. Where did you obtain most of your views/knowledge of AI? We guessed that most people would probably select “Movies/books” for their answer due to the large amounts of artificial intelligence in science fiction films and literature (and television also), and the assumption that the majority of people taking the survey were not AI developers, students, teachers, etc.

Q9. What is your profession/occupation or field of study? Q10. What is your age? Q11. What is your gender? These questions were asked purely for statistical and comparison purposes. We simply wished to see if age, gender, or field of experience/education had any impact on the way that people answered the questions.

Now let us take this a step farther and consider some of these questions in combination. It is our belief that people who find AI a threat to society have gained most of their knowledge about it from movies and books. Similarly, people who gained most of their information from class or work find that it is ultimately a way to improve life. The reasoning behind this is that a very good sci-fi or action thriller can be made by having a robot that revolts against its creators. Films such as *The Matrix* and *Terminator* exemplify this very well. People who have learned about it in class or work with it in the field a far more likely to see the great the things about it, and all the possibilities it holds.

Women had to fight for their rights, so would they be more likely to want to extend them to robots? We thought so. We also figured that more women may see robots as creatures than men would.

We also thought that a person’s profession/occupation would determine how they viewed AI and whether or not they thought that it would become a problem in companies.

A janitor would view robots as more of a threat than a human resources professional, for example.

7.2 Survey Results

The results of the survey conducted can be found in Appendix C. The survey was taken by 197 people. Out of this group 70% were male and 25% were female (the other 5% were unknown). More than half of the people who took the survey were in their twenties. People in their thirties and forties made up another quarter of the people and the remaining quarter consisted of older people and teens. A good portion of the people, 39%, dealt with computer science. 20% of the people were engineers. The rest were scattered between various occupations or fields of study.

In asking when they think true AI will be achieved, 38% of the people said in 25 years and 43% said in this century, with the rest either saying later or never. But when asked whether machines will be able to think on their own, half of the people said yes and half said no. Almost everyone thought that AI will find a place in our everyday lives in the future and two thirds of the people felt that there will be a problem with AI taking over people's jobs. The remaining third believed that there will not be a problem. Only 9% of the people see AI as a threat to society and 69% see it as a way to improve life. Although 75% of the people felt that intelligent machines shouldn't be given similar rights as humans.

About a quarter of the people who took the survey obtained most of their knowledge from movies. Another quarter obtained it from the news. The rest learned from work, class, or a hobby. More than half of the people said that they interact with AI

daily or once a week. 17% of the survey takers said that they have never interacted with any form of AI before.

7.3 Survey Interpretations

(For charts referred to in the next section, see Appendix C)

Interpretations of these graphs are important because it breaks down the results from the survey and shows you specifically how different groups of people view AI and what trends and similarities exist between these groups. Through some interpreting, assuming, and thinking, one can easily see that thoughts and views almost always depend, change, or even exist because of differences in opinion. Also, through trends or experiences in gender, age, and occupation, a cornucopia of observations could be made. Someone who has worked with computers for years and has a computer-related job most definitely believes that AI is beneficial and can only improve life. One who holds such a position in life might think highly of or take pride in something involving their field. Where as someone who is much older and does not have a computer might believe that they are useless contraptions that could someday overthrow society. Hence there is a need to analyze the survey results and draw some comparisons between the responses.

Chart #12 shows what percentage of people answered yes, no or did not answer the question about the future of intelligent machines and whether or not they would hold the same rights as a human would in comparison to the results about when AI will be fully achieved. An interpretation of this shows that about the same percentage of people who answered yes or no agreed as to when the development of artificial intelligence would be realized. More than half of the people surveyed, however, felt that AI would

never be achieved. The graph also shows that the majority of people who did not answer yes or no to the question on machine rights, mostly thought that AI would become a reality sometime this century.

Chart #13 compares professional occupation or field of study and when AI might be achieved. The bars show that more than half of the people who took this survey had some sort of computer or engineering occupation. Although, their thoughts on when AI would be complete was somewhat evenly dispersed between 25 years and this century (about 47/48 people), as well as later and never (about 10/12 people). This goes for the rest of the occupations as well. Close to an equal amount of people from one profession thought either AI would be achieved within 25 to 100 years, and a lower amount of people in that profession about the same for later or never.

Next, Chart #14 asks about the future of AI. Where does the person taking the survey see AI going: as a threat, as a way to improve, indifferent, or unanswered? This is in comparison to if the person also thought that robots would someday replace humans. More than half of the people surveyed believed that yes, someday robots would replace humans in some companies or at jobs. The majority of people thought that this would be an improvement to life, where others believed robots would be a threat to society. Also, there were more people who were either indifferent or did not answer than people who thought it would be a threat. An even part of the unanswered answers to how they saw AI left the other question unanswered as well.

Another graph, Chart #15, shows a comparison between age and if machines will be able to think on their own. Only one person did not tell their age and did not answer the question, but still, few people who answered yes or no, did not list their age. More

than half of the people who took this survey were between the ages of 20 and 29. More than half of the people who took the survey also answered yes to whether they feel will machines be able to think on their own. It seems that the older the survey takers were, the higher percentage of them answered no to the question of machine thinking ability.

Comparisons and contrasts of how people have learned about the different aspects of AI as well as where they see it going are shown in Chart #16. More than half of the people that believe AI will improve life split up into even groups between the various knowledge sources. Very few people think that AI will be a threat, four of which think so from a class or from work, where the other thirteen learned from movies, books, or the news. A good amount of people are also indifferent, which does not really hurt or help, because of the huge difference between thinking AI is a threat or a helpful tool.

Chart #17 compares professions and how those people see AI. About 1/3 of the people who took this survey do work involving computers and believe that AI will be a way to improve their lifestyles. Another large chunk of people who believed that AI would be an improvement, were the engineers. Few people from each occupation think that AI is a threat to society, and even more people with various occupations are indifferent to the question.

Another graph, Chart #18, shows the comparison between age and how they see AI. It seems that the majority of young people between the ages of 20 to 29 believe artificial intelligence will improve living. Few younger people and many older people who took this survey were indifferent to how they saw AI. Also, very young people saw artificial intelligence as a threat to society more so than very old folk.

The next graph compares how the different genders see artificial intelligence, shown in Chart #19. More than half of the people who took the test were male and thought that AI would improve living. Also, more males were indifferent to the question rather than thinking that AI would become a threat. Most of the females thought that AI would improve life, and again, more women were indifferent over thinking that AI could be a threat. Finally, the people who did not specify a gender, more of them (2) think AI will improve life and the other (1) thinks AI is a threat.

The final graph shows the comparison of age and how often the person interacts with AI in Chart #20. The majority of people surveyed said that they interact with AI on a daily basis, about 59% of the people were between the ages of 20 and 29. Each age group and distribution of people in each category were all dispersed somewhat regressively between the amount of time, decreasing from daily to once a year. Although, there were a good amount of people surveyed, 17%, that feel they have never interacted with any form of AI.

7.4 Summary

The survey was taken mostly by men and people in their twenties. Also, more than 50% of the people who took out survey were computer scientists or engineers. Almost all the people thought that AI will be truly achieved. The two main places people learned about AI was from either news or works of fiction (movies and books). This coincides with our predictions of these results. The majority of people felt that AI will improve life, which we predicted. We also thought that most people would fear that AI will take away people's jobs, and indeed this was the case. Contrary to what we

expected, there was a large age distribution when asked how often people felt they interacted with some form. What's more unexpected is that the majority of people in their 40s said that they interacted with some form of AI daily. Mainly people feel that the human race will benefit from AI in aiding and carrying out everyday tasks. The only problem our survey takers really seem to think AI could hold is the danger of people losing their jobs to intelligent machines. This could probably be accounted for by the source of information the people who took the survey are basing their knowledge on. The media doesn't largely publicize obvious or helpful forms of AI, but rather those that are "threats" or not easily recognizable as artificial intelligence. Thus AI's beneficial side isn't always seen nor are its limitations.

8. Moral and Ethical Questions

The area of artificial intelligence is riddled with ambiguities and misinterpretations. Many people feel that artificial intelligence will never be achieved while others feel that it is already a reality. Such discrepancies can only come about if there exists an inadequate definition of what artificial intelligence is. One possible reason for the lack of a definition may stem from the lack of a unified set of goals. What exactly is AI trying to achieve? One thing we will attempt to do in the following section is come up with a definition of AI that is based closely on the definition of the individual words and the origin of the term itself. Also, we will try to use this definition to clearly state what the general goal of artificial intelligence is currently and where it might be headed.

Other points of interest in the proceeding section include what responsibilities we have to the intelligent machines we create. If they truly are intelligent do they deserve to be treated as an intelligent being is treated? Also, what responsibility do we have for ourselves, to protect human interests? This subsequently raises several philosophical questions at the root of the problem of many areas of AI. Can machines be social, ethical, and even spiritual and what defines each of these terms?

8.1 Definitions and Goals

As stated above, several problems pertinent to the field of artificial intelligence stem from the lack of a single, clarified definition of AI itself. For example, look up artificial intelligence in any dictionary or encyclopedia and see the number of varying definitions for the field. (For a summary of the most common definitions see the website <http://www.wpi.edu/~kerrym/ai>). Common entries range from claiming that AI is the

study of how to make machines act like humans, to how to make computers solve problems, to how to make them “do the things they do in the movies.” No wonder there is so much confusion surrounding AI, even people who are familiar with the ideas of it still can’t settle on one definition. One might say, “What’s the big deal? Just pick a definition already.” Unfortunately the root of the problem does not lie with people arguing over something as trivial as a simple definition. The problem stems from the largest discrepancy in the AI field: what is the ultimate goal(s) of AI. If the only goal of AI is to achieve superhuman calculation prowess then it has already been achieved. If the ultimate goal of AI is to become as human as we are, then it may never be fully realized. Of course those are just the two extremes of the spectrum. These are the real issues that people are divided on and due to these varying goals a single, universal definition of AI does not exist.

Let us start with the root of the problem. What is the goal of AI? A commonly cited goal of AI is developing computers or machines capable of calculating extremely complex problems such as guided missiles, proving mathematical theorems, pattern and voice recognition, weapons and aircraft design, etc. If this were the only goal of AI, then I feel we have all but achieved it at this point. The only thing that AI would have left to work on would be speed optimization and time efficiency issues since we have many computers and programs to do just this, calculate at fast rates, large amounts of information. This can’t be the only goal of AI. There must be more.

One common perception is that artificial intelligence is the study of how to make computers “like us” or similar to humans. In this respect there is much ambiguity. How far do you take the “similarity”? Is it only so far as intelligence is concerned, as the name

would imply, or are we trying to create replicas of ourselves that not only act and think like us but also look and move like us as well? At this time we are much closer to replicating the capabilities of a human brain than we are to incorporating that into a bio-mechanical being. The idea of “cyborgs” or “androids” is a bit too far-fetched to earnestly research with our current technology. Also, artificial intelligence insinuates that you are trying to create just that, artificially intelligent “machines”. Nowhere does the name imply artificial life. Therefore the idea of making machines “like us” is more or less limited to intelligent abilities and nothing more.

Surely though we do not wish to only create a “fake brain”. What would the point be? There are over 6 billion people on the planet, each with their own unique brain. What purpose would creating an artificial replica of a brain serve? Well, one possible use would be in the psychology field. Recreating a mental process would be exponentially helpful in the field of studying psychological problems and how they manifest. Worlds of knowledge and understanding await researches if we were to replicate human thought artificially. We would have “cracked the code” so to speak. However, with the possible exception of some biological studies, there is not much more that this fake brain could provide us with. It would be a great achievement indeed, but ultimately useless. Along with recreating human thought you also recreate human fallacies. There are many things humans can’t do. Surely you would not want your computer to sit and think about the answer to a complex math problem as it counts on its artificial fingers. Similarly, you would not be happy if your new computerized stock broker made a foolish market decision and ended up losing millions of your assets. Humans are capable of bias and prejudice. They are able to be swayed by emotions and hormones and are susceptible to

chemical imbalances and mental handicaps. "Artificial Stupidity (AS) may be defined as the attempt by computer scientists to create computer programs capable of causing problems of a type normally associated with human thought." - Wallace Marshal

There are some things that computers inherently do better than people. For this reason, we conclude the true goal of artificial intelligence is to create super-intelligent computerized "minds" that are capable of creativity, informed decision making, flawless memory, and extremely fast processing times. The result would be a human-like mental entity that can out perform the human mind in every way and which shares all of our advantages as humans without any of our drawbacks. But what happens when this goal is met? This goal does not, with today's trends in technology, seem unattainable. Where does AI go from there? Perhaps then an incorporation of this technology and advances made in biotechnologies could be made. Memory enhancements, audio augmentations, "intelligent prosthetics", or even full-fledged androids, the list continues. Unfortunately, we are not studying time travel, so there is no way to tell what the future of AI will hold.

Based on these goals, perhaps the best way to define artificial intelligence would be to base it on the definitions of the individual words, namely artificial and intelligent. According to <http://www.dictionary.com> the word artificial is defined as "made by humans, produced rather than natural" and also as "made in imitation of something natural, simulated". Furthermore, intelligence is defined as "the capacity to acquire and apply knowledge". Artificial intelligence then, could be described as the study of how to produce an unnatural entity capable of acquiring and applying knowledge. This definition effectively takes the middle ground of the two sides of the argument surrounding AI. It does not insist that AI is only about fast calculations since

even a very fast computer or calculator is still given the information it uses and does not acquire it on its own however it also does not try to claim that something artificially intelligent **must** be a conscious, self-aware entity that mimics human thought in every way. This broad definition accommodates nearly all schools of thought in AI in some fashion. All of the different “levels” of AI, such as simple calculators to guided missiles to androids with emotions, are all defined under the same term since this definition also allows for the expansion of the AI field without a necessary change of definition of what the field ultimately studies.

8.2 Philosophical Questions

What defines thinking? In the previous section we defined intelligence so it may seem as if we are being redundant, but are we? Is there a difference between intelligence and thinking? There is no doubt that a machine can be at least grossly intelligent. Our 99 cent pocket calculators can vastly out perform us in mathematical operations and today’s supercomputers can beat us in chess, so intelligence is not something machines are lacking however no one would really argue that these machines are “thinking”. Thinking is defined as “a way of reasoning or judgment” it can also be defined as “pondering”. Surely a simple calculator is not passing judgment with the numbers you give it nor was Deep Blue reasoning against Gary Kasparov during their famous chess match. To this extent machines do not yet display signs of truly meaningful thought beyond what their algorithms allow. However this is not to say that machines will always be incapable of thinking. When one actually sits and asks themselves how they are able to think, no one can really answer. Descartes’ famous quip “I think, therefore I am” fails miserably if it

can not be shown how this thinking is accomplished. Who is to say that the way humans “think” is different than anything a machine could ever do when we can’t even understand the process that gives us this privilege? It is argued that machines are merely following lines of code and operate only by an algorithm. This rigid structure renders them incapable of true thought since creativity and ingenuity by definition require thinking outside of what we already know. How do we know, however, that how humans think is not also by an algorithm? Humans could very well follow an algorithmic approach not too dissimilar from that of a computer. The human mind is, for all intents and purposes, a very advanced and dynamic piece of software. With the advances made in technology in just this last century, advances that would have been considered impossible to someone in the 19th century, the truth is we can’t predict what new programming approaches or technologies will be in store. It is just as ignorant to say that machines will never be able to think as it is to say that they definitely will. For the purposes of this paper however, we will accept the possibility that machines could one day think. This however does not answer anything but instead only opens up the door to several other philosophical questions.

AI machines are being created, as discussed earlier, in our own likeness. Their thought processes and abilities are molded after humans. If artificial intelligence truly does succeed in giving the ability to think and reason to an artificial being such as a machine, are they really capable of meaningful thought and ingenuity? After all, they are just mocking human abilities. Does mimicking human thought constitute human thought or is it just a good “parlor trick”? When a dog is taught a trick or when a parrot mimics a person’s voice they are not being creative. If a computer mimics us then why should they

be considered any different than a parrot? The difference however is that a parrot is blindly recreating the sounds it hears without having any idea of what it means. It may not be thinking but it is “talking.” A computer that can think like a human, whether or not it is only mimicking the human, is still thinking like a human. The mockery of humans is merely a technicality. Actually, it is a human drawback that we are only able to give machines our abilities since we ourselves are incapable of giving them otherwise. Once machines are able to think like a human and be creative like a human they may be able to better themselves in ways we could never have seen due to their increased speed and efficiency.

Another pending question is if credit for advancements or jobs performed by a machine should be given to the machine’s creator or the machine itself. Isaac Newton, Albert Einstein, and Thomas Edison are all famous names in history books for their discoveries and scientific achievements. We do not, however, give credit to their parents for their achievements, and rightfully so. They were separate people with distinct personalities and skills. Similarly, we don’t hear much about Michelangelo’s father or Galileo’s mother; the creators are not given credit for their offspring’s work. With machines the question becomes more complex. A machine is not merely the blind biological product of a couple’s intimacy. Work, effort, and ingenuity does not go into the birth of a child, therefore when the child turns out to be a gifted thinker, the child is given the respect. An intelligent machine represents the hard work, devotion, and creativity of a person or group of people. Whatever the machine is capable of has been given to it by its creator and programmer. If a machine becomes creative, however, and starts doing things it wasn’t originally capable of doing or achieving, even according to

programmer's plans, does this artificial evolution belong to the machine or its creator? We think a good answer to this question is a combination of both. For the first generation of intelligent machines that are produced, credit should be given to the programmers and designers of the technology. As generations of machines pass by and quite possibly, machines begin to make new machines of their own, any further advancement made devoid of human intervention should be credited to the machine responsible for it. For example, Deep Blue should not be given sole credit for beating Kasparov in Chess, its programmer is deserving of that title. Deep Blue's hypothetical great grandson, Perfect Blue, capable of human thought and itself the product of machines, is entitled to respect.

The discussion of thinking machines will now switch gears slightly to the topic of emotion and feeling. If a machine is made to be capable of thought could it also be made capable of emotional response? Can machines one day laugh and cry, get sad, be vengeful? Maintaining the hypothetical, let's assume that it is possible for this to happen. Even today it seems possible to recreate emotion at the most rudimentary of levels, such as with virtual pets. Does this qualify the machine as being human-minded or is there some intrinsic value to the human mind, some intangible virtue that cannot be distinguished nor refuted? As is the case with most philosophical quandaries, this question raises a number of other questions long debated by philosophers. Should it make a difference, philosophically if the human thought process is emulated as opposed to simulated? As described earlier, emulation centers on the end result while simulation focuses on how the result is achieved. If machines are capable of serving the same purpose as a human but do not reach these ends in the same fashion we do, does that make a philosophical difference?

The question at the root of the discussion boils down to whether or not you believe that the mind is separate from the brain. Several religious questions surface from this which we will not get into in much depth here. Briefly, in many organized religions the mind is synonymous with the soul, a separate dimension to the human being. Separate from the biological body the mind is a part of us that lives on after our earthly ties die

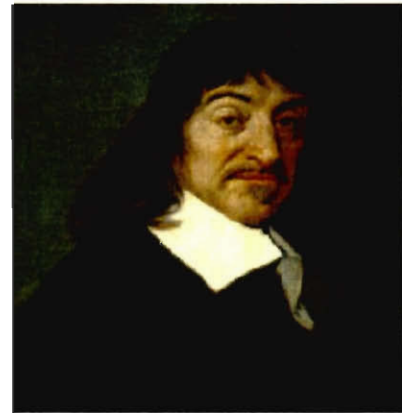


Figure 8.2.1 – Renè Descartes (1596 – 1650)

off. This ambiguity in what exactly the mind is makes discussion of it difficult and prevents any useful arguments to be made. The mind is immaterial, a disembodied intellect, case and point. This was the view held by aforementioned philosopher Rene Descartes. The theory is known as “Dualism of Mind and Body”. The real world can lie; it is subjective to how our bodies interpret our experiences. For all we know it could be a dream, but what separates ourselves from the falsehoods of reality and from our bias of perception from person to person are our minds, according Descartes. Another school of thought is that the mind is not separate from the body. A mind is symbiotically connected to the body. A different mind in the same body would be a different person just as much as the same mind in a different body would be a different person. This view of mind and body unity makes the word “mind” nearly interchangeable with the word “brain”. One supporter of this view was the philosopher Martin Heidegger. Heidegger believed we were not immortal souls but finite historical existences. The truth behind the connection between the mind and the body may never be fully known or understood. Until that time we can only speculate and the answer to the question if a machine is

capable of achieving the type of spirituality a human possesses is left to the discretion of the reader's own personal beliefs.

Lastly, cyborgs present an interesting dilemma. Part machine and part biological entity, in some respects cyborgs already exist. In the not too distant future AI technologies and biological advancements could potentially merge together to form kind



Figure 8.2.2 – Martin Heidegger (1889 – 1976)

of “super” humans. Memory enhancement devices, artificial limbs, internal hearing aids and visual enhancements allowing for the ability to hear and see things beyond the scope of the natural human; the list continues. The idea of microchips being installed in people is not a new concept. In fact there was talk of installing such a device in children so that parents could keep track of where their kids were at all times. Not too many people

argue that a person who just gets a hearing device is no longer the same person but what if it didn't stop there. What if after the hearing augmenter the hypothetical person had a memory enhancement device or perhaps an extra “processor” installed so that he/she could think faster and more efficiently? Maybe he/she would get their limbs replaced with synthetic ones that were stronger and more reliable. At the extreme a person could theoretically download their neural structure into a synthetic neural network creating a duplicate of their mind for backup or us in a different way. (Kurzweil, 1999) Is the backup copy also the person? How about an intelligent robot using the person's mind? Is the bionic person still the same person just faster and more powerful? Answers to these questions are not

something we can honestly give. Philosophers themselves are still unsure how to deal with this interesting situation artificial intelligence brings about.

8.3 Social Aspects

Can a machine be a social entity? Intelligence is impressive and even the ability to think and reason is amazing, but what about social interaction. Stemming from the discussion of feelings and emotion, sociability could be looked at as the readiness for one to form groups and associations with others. The basis for forming social groups is to satisfy a need. The primary need that groups satisfy is support; physical and emotional. Comforting oneself while in the presence of others is a common way to relieve stress and be with those that share similar experiences. Feelings of isolation are usually caused by a lack of social interaction. Now it would not be unreasonable to assume that a computer could be at least programmed to simulate sociability so as to make it easier and more natural for a human to interact with it. Whether or not the computer would actually be feeling true emotions and displaying genuine affability is something that may never be known to anyone else besides the computer. As stated in a quote from *2001: A Space Odyssey* in response to a question whether or not the astronauts felt that HAL had genuine emotions, "Well, he acts like he has genuine emotions. Um...of course he's programmed that way...to make it easier for us to talk to him. But as to whether or not he has real feelings is something I don't think anyone can truthfully answer." The real question then may not be if machines can be social with us but if we can be social with machines.

Once in the business and social sector however, regardless of whether machines are truly capable of being sociable, should there be a limit to how much machines are allowed to interact with humans directly? Supposing machines are programmed to be reliable and pleasant to work with, why should there be a limit to how much exposure exists between machines and humans outside of the research and scientific sectors? This however leads to another question which will be addressed formally in the Risks and Perils of AI section. How much control should be given to machines and what should remain under human grasp?

Affecting a machine's ability to be truly social, emotional, or perhaps even spiritual could be its level of consciousness. An indicator of meaningful thought is self-awareness. If a machine is capable of becoming truly self-aware; more than simply knowing it exists but knowing that it knows it exists, then machines could in fact be on the road to consciousness. True consciousness is not just acting out on self-need. Plants move to the direction of light and some of them are even capable of defense mechanisms yet not a one of them is "conscious." For a machine to really display consciousness it would have to display free will, the ability to make decisions based on a balance between self preservation and selflessness. A robot programmed to only defend itself or one programmed to only protect others is not a conscious being, it must necessarily be capable of both at its own discretion. It can be argued that the same theory could apply to humans as well.

8.4 Responsibility

In the 1960s the topic of civil rights became a prominent issue in America. Feminism, animal rights, equality of rights amongst the races; civil liberties came under intense scrutiny during this era. Modern liberalist movements share many of these roots in common. Environmental protection and abortion, for example, use many of the same arguments proposed by animal rights and feminist activists. Why should these basic principles of equality end with the natural world? In the future, when machines become intelligent entities and important parts of our everyday lives, the issue of civil liberties will once more be brought to the forefront of our culture.

Are we morally obligated to the machines we create? Before we answer that question however, it is necessary to ask what makes something worthy of moral consideration to begin with. Why is a tree, an animal, or even a human deserving of morality? An argument used by ecologists and animal rights activists for a long time is that while lacking in the level of intellect possessed by humans, plant and animal life is capable of response to our actions. We can affect their health and ability to survive. In short, an animal can experience pain and can die. A tree, while unable to experience “pain” in the sense that humans know it can endure harsh treatment and eventual death. The ability of something to suffer then is our connection to the natural world, the one thing we share in common. Still though, why should a human be worthy of our time and consideration? The simple and trivial answer is “just because”. That is the common response you will get from a person when asked why they feel someone should be treated in an ethically sound manner. It seems like the answer is just common sense, however very few people can actually give an intelligent response. One famous argument is that

we treat others the way we wish to be treated. This however does not explain much. It assumes that we inherently trust others to also abide by this convention, an obvious impossibility given the amount of crime and suffering in the world. The answer lies in virtue. We seek to be virtuous. For most people it is virtuous to be a “good person”, to treat others the way they would expect to be treated. For some, are mentally ill, they feel they are virtuous by killing or committing crime. In their own twisted fashion, they are adhering to their own code of ethics. Whatever the impetus, virtue is a driving factor behind morality. Therefore, if it is virtuous to treat humans with moral consideration because they are able to suffer as you can then it is also virtuous to treat animals with at least a small degree of moral consideration for the same reason. Furthermore, through a similar argument, plant-life and Mother Nature is able to suffer and die, therefore they too are deserving of morality.

This doesn't mean you should become a “Save the Rainforest” activist as soon as possible. It depends on the definition of what you consider to be virtuous behavior. Some people feel that only intelligent beings, such as humans, have morality. Animals are just beasts, drones without the intelligence or capacity for moral thought. In some sense they could be viewed as machines. One of the problems animal rights activists face is trying to overcome the stigma that animals are just machines we have dominion over. A dog experiences pain, a chimp can use tools and sign language, and even lesser species such as snakes and other reptiles exhibit instinct and abilities far beyond that of humans. To anyone who has ever owned a pet they will most likely find it difficult to honestly view their companion as a machine and nothing more. Similarly, artificially intelligent machines could be viewed the same way. Whether or not they ever achieve true

emotions or consciousness or even their own sense of morality, it does not affect our obligation to be moral towards them. A machine that is intelligent and capable of human thought, can converse with you, perhaps one that can even act like a pet, is this not similar to a dog? Furthermore, like the snake, machines will exhibit abilities humans cannot even come close to possessing. An intelligent machine is also fully capable of “dying”. Disassembling, destroying, draining, deleting; whatever the method a machine is not impervious to termination. Also, machines could be programmed in the future with the ability to feel pain in a similar fashion to humans, affecting their ability to perform. In fact, the ability to feel pain could be the key to unlocking other higher level mental capabilities such as morality and genuine emotional response.

Returning to our previous discussion on civil liberties then, it does not seem farfetched to think that advanced, well informed machines should be kept from having some of the basic rights that humans also possess. If a machine could evaluate more objectively and more precisely what would be the best direction for society to head in or which candidate’s platform is more beneficial to the nation, there should be no reason to prevent these machines from having the right to vote. In fact, perhaps it would even be more effective to allow machines to control some portions of the society and government. Human errors and biases, campaign contributions, corrupt politics, personal interests; all of these hinder the governmental process and only bring down society as a whole. The ability to eliminate these inefficiencies could prove invaluable to the advancement of world culture. Of course machines would also be held accountable for their actions by the same set of laws that govern human behavior. Sound farfetched?

Now that we've examined the possible ways we are responsible for the machines we create, we should examine what responsibility we have to ourselves. If the advancement of machines results in the displacement or suffering of our own species we are morally obligated to ourselves to prevent such a thing from happening or allowing to progress beyond our control. There has always been a fine line between technology and terror. Nuclear power has done much to help alleviate power generation issues and fossil fuel consumption greatly. Unfortunately it is also gave birth to one of the most destructive and terrorizing advancements of the modern age, the nuclear bomb. Governmental regulation cannot always be trusted to prevent such atrocities from becoming reality and the disguise "in the name of science" can only take you so far. The new millennium, and in fact, the new century will contain many interesting advancements, some of which we can't even begin to predict or understand. Human control and interaction with technology will become crucial to the survival of our species. With this in mind, let us take a more in depth look at some of the potential risks and perils presented by artificial intelligence in Chapter 9.

8.5 Summary

Some of the more important thoughts discussed in this chapter dealt with the moral and ethical questions raised by artificial intelligence. Firstly, it was necessary to set a basic goal and definition for AI so we could then use this in our discussions of the moral issues the field raises. Philosophically speaking the premise of a machine being able to think seems to contradict many preconceived notions of the human mind. If a machine is able to truly think than humanity's claim to the anthropocentric throne we've

held for the last several thousand years is seriously jeopardized. Mankind has a tradition of thinking that it is utterly unique and special from the rest of the world. Intelligent machines pose a possible threat to that view, forcing us to reevaluate our own minds and morality. Also, it is interesting to speculate what kind of impact on society intelligent machines will have along with what kind of impact we will have on intelligent machines. Potentially we could find ourselves amidst a new civil liberties dispute, this time between man and machine.

9. Risks and Perils of AI

Many of the questions raised in the previous chapter lead one to wonder if humans should be worried about further advancements in artificial intelligence. Not everybody thinks like Ray Kurzweil. Welcoming a utopian society where the lines between man and machine blur until humans eventually yield to the control of machines



Figure 9.1.1 – Ray Kurzweil

may sound like a good idea to him but for many people that is a frightening prospect. Is it truly possible though? If we are to believe movies and television and science fiction, the answer is an emphatic “Yes!” Those of us more grounded to reality than your average conspiracy theorist and science fiction fan however are not as quick to buy into the mass hysteria. The idea itself seems wholly unbelievable. To answer this question we need to first address some other issues. How dependable are machines? How dependable are we? What will human-computer compatibility be like in the future? What is the Luddite movement and how does it relate to AI? Lastly, is AI really a threat?

9.1 Reliability

Machine dependability is an important concern in many current applications. Cars, computer chips, automated factories, roller coasters, operating systems other than Microsoft Windows; they are all designed to be dependable and reliable machines. The over reliability of these machines however, depends greatly on the reliability of their creators and designers. In the future the same principle will still apply. Intelligent

machines will be made to be dependable and reliable but the end result will depend greatly on the inventors of the technology. So the question of machine dependability really comes down to a question of our own reliability. Some people are very dependable while others aren't quite as trustworthy. In a similar fashion some future intelligent machines could be very reliable and some could be unreliable. Furthermore, if machines are indeed modeled after the human thought process it is conceivable that, just as in humans, corruption of machines could become an issue. If the human brain is modeled accurately enough, there is not much to prevent a machine from acting based on its own interests or becoming biased as a human does.

According to Isaac Asimov's three laws of robotics, it could be possible to design robots that are not capable of malfunctioning and would be completely reliable. The three laws are as follows:

1. A robot may not injure a human being or, through inaction, allow a human being to come to harm.
2. A robot must obey orders given by human beings except where such orders would conflict with the first law.
3. A robot must protect its own existence as long as such protection does not conflict with the first or second law. – (Asimov, *I, Robot*)

This of course was not meant to be the absolute guideline for robots to be designed by since Asimov was writer of fiction more than philosophical or technical thought. However many researchers have used his suggested laws as the basis for advancements in robotics. These laws are similar to the primary objectives programmed into the machine

from *Robocop* discussed earlier. There exists a few drawbacks with this theory however. First off, these “laws” are very constrictive. While theoretically sound and perhaps the safest for humanity they also prevent a robot from achieving true consciousness and free will as described in the previous chapter. Part of what could make emotion possible lies within the ability of machines to make their own decisions. A famous line from another one of Stanley Kubrick’s movies, *A Clockwork Orange* (1971), is “If a man cannot choose, he ceases to be a man.” Similarly, the machines described above are devoid of choice. They have to act in a predetermined fashion. Part of what true intelligence and thought entails is the ability to make educated choices based on information. Sometimes these choices sacrifice the good of one for the betterment of the whole. There is no room for sacrifice other than self-sacrifice by the robot in the aforementioned set of laws.



Figure 9.1.2 – Isaac Asimov (1920 – 1992)

Another problem with the laws is the inherent trust they place in the designers of the robots. As stated above, the robots will only be as dependable as their creators are (at least initially). If almost all AI advancements were made in accordance with the three laws of robotics it would become a relatively simple task for someone to deliberately disobey the rules and create a robot not governed by the same set of guidelines. This robot would be at a great advantage since it would not be hindered by intertwined laws preventing it from harming humans or having to obey them. These laws therefore do not exclude the possibility of robot revolt or corruption since it again relies heavily on human dependability.

9.2 Human-Computer Compatibility

Human-computer compatibility is another potential problem AI poses. As discussed previously, humans have a history of fearing those things that are new or different. For many centuries African Americans were feared because they were “different.” In more recent times homosexuality has become a hot topic. Fear of the unknown or misunderstood is a powerful catalyst for controversy, dispute, and even war. Couple this with the Frankenstein Complex (man’s tendency to fear technology) and you have one volatile crisis in the brew. Not only will AI become an important and integrated part of society, but if intelligent machines enter the public sector and social forum, the fear of what they could be capable of would become a widespread epidemic. Most of the time this type of fear is unwarrantedly over-exaggerated however that is not to say there is no reason to be concerned. Stated in the previous chapter, civil liberties for intelligent machines could become an issue. Much like previous and current issues such as racism, gender bias, and homophobia, the civil rights debate as applied with intelligent machines could spark unprecedented amounts of protest and turmoil. Whereas a gay man or an African American is still undoubtedly a human being, worthy of moral consideration in some form, an intelligent machine faces a more difficult situation. People accepting machines as equals or at least as entities deserving of rights is not something we predict will happen smoothly or quietly, if it happens at all. If humans and machines cannot “get along” and coexist peacefully, many problems lie ahead.

Human-computer compatibility can be viewed from a different perspective, perhaps in a more literal sense of the term. A common phrase heard in computer technology today is “user-friendliness”. A computer or machine that is hard to operate or

work with is often considered bothersome and inferior. For intelligent computers of the future the same will hold true. Compatibility between humans and machines must be made a priority. If a human cannot comfortably interact with a machine, both the human and the machine will be at a disadvantage. By interaction we mean several different things. On a purely social level, interaction could mean ease of information sharing or conversing. On a more technical level, these intelligent machines would have to also be “user friendly”, that is easy to use and work with. For the AI field to really take off, mass market acceptance and appeal is a necessary requirement. People will not want to be hassled with learning complex machines or unnecessary technical requirements in order to get the functionality they desire.

What is a Luddite? The definition of a Luddite is as follows. “One of a group of early-nineteenth-century English workmen who destroyed labor-saving machinery in protest. The Luddites were the first organized movement to oppose the mechanized



Figure 9.2.1 – Artist’s Sketch of Theodore Kaczynski (a.k.a. The Unabomber)

technology of the Industrial Revolution. Today, the Luddites are a symbol of opposition to technology.” – (Kurzweil, *The Age of Spiritual Machines*, p. 305) The connection the Luddites (or Neoluddites) have to artificial intelligence is actually rather disturbing. Theodore Kaczynski (a.k.a. the Unabomber) was a madman responsible for the death of several people and the injury of many more through mail bombs he would send to specific people, usually to professors at colleges and universities

over nearly a two decade period. He published a manifesto concerned with the Luddite movement shortly before his apprehension and imprisonment. In the essay, he

documents the evils and atrocities caused since the Industrial Revolution as a result of technology. Specifically, he makes it a point to target AI, as the root of much trouble in the future. The frightening part is that he actually makes some sense. If the paper is read without the knowledge of who its author is, it surprisingly holds literary merit. He proposes various scenarios of how the future could pan out and predicts what would happen for each situation. Overall, his belief is that technological advancement should not be allowed to progress further. If Kaczynski's manifesto isn't a testimony to the volatile and potentially perilous situation concerning human-computer compatibility and the potential for animosity to exist, nothing is. For the full text of Kaczynski's manifesto visit this website, <http://www.soci.niu.edu/~critcrim/uni/uni.txt>.

9.3 Is AI a Threat?

Artificial intelligence is a slippery slope. At first advancements could be rather small, such as artificially intelligent stock machines or AI automated vehicles that can taxi you from place to place. While these are impressive achievements they don't really pose much of a threat. However there is not much of a technological leap between AI taxis and stock quote machines to full fledged artificial stock brokers in control of the market. Allowing small advances in the field could potentially open the flood gates to many unwanted side-effects such as job displacement. The AI field, like so many areas of technology, is advancing too fast for the rest of society to keep up. Much of the population is just getting used to the Internet and already there is something new on the horizon. The human genome has been fully decoded yet many people still don't know what that means or represents. The cloning of a human seems almost inevitable now

when just a few years ago it was considered impressive that a frog had been successfully cloned. There are no laws to regulate these scientific achievements yet since our law system moves exponentially slower than our research. While philosophers still debate over the mind and body connection, animal rights activists fight to save the whales, and record companies argue about how to handle royalty fees and piracy issues with mp3's over the Internet, a new race of beings is quickly emerging faster than we know how to deal with them. If this paper has shown anything it has shown that AI only raises more questions than it does provide answers.

Is AI a threat? Well the short answer is yes. As with all technology, the potential for disaster exists. Therefore artificial intelligence does pose as a potential threat although we are more inclined to agree with the majority of our survey takers that AI will ultimately be a way to improve life. However, there is no denying the catastrophic possibilities it could have. Unmanned warheads and nuclear weapons in the control of machines is a scary thought but its real life implementation is not far off. Concerns about the reliability of such a loss of manual control and especially job displacement are not unwarranted. Call it blind faith in humanity, however, that we will be able to find a way to prevent such a thing from happening and that we have an effective way of dealing with it in case it does.

9.4 Summary

Could humans become obsolete? If computers have the potential to become so much more advanced and efficient than the biological mind is it not possible for humans to suffer from a real-life version of the Frankenstein Complex. At the rate technology is

progressing today, will we have enough time to consider the consequences of what we create before it's too late? What would prevent our creations from revolting against us or treating us as we treat animals now? We could potentially become victims of the modern age, lesser beings unable to stop what we've created. Science fiction isn't always wrong. Just read some of Jules Verne's novels and see all of the predictions the gifted writer made that have now become reality. Machines of the future might not have a need for human intervention. We would become a bothersome nuisance, a civilization that does nothing but hinder the progress of the intelligent machines and use up valuable natural resources. Are we destined to end up as the battery power for the machines of the future, plugged into a virtual reality as in *The Matrix* or fighting for our lives against killing machine androids like in *The Terminator*? While those stories definitely rely on exaggeration and embellishment of reality for entertainment purposes, the truth may not be too far off from such a scenario. Of course this is not meant to scare the reader. A line from *The Terminator* sums up how uncertain the future really is, "The future is not set. There is no fate but what we make for ourselves." It is entirely possible, as Kurzweil suggests (and looks forward to), that the machines could act like grateful children towards us; superior to their elders but respectful of them for giving them life. Before we pass into the history books as a species that once was, we would breathe our last evolutionary breaths as an elderly parent in a nursing home; made comfortable and surrounded by all the help we need before it is finally our time to leave and hand over control to our artificial children. Of course, anything is possible. One thing's for sure, it's going to be a fun ride.

10. Conclusions

Technology today is evolving at such a fast rate that people can't follow everything that's is going on in the industry and may not understand it. It is very important that the public know about the rising technology so that they behave in the appropriate manner towards it.

Artificial intelligence has a very high probability of being misunderstood due to bad press and general ignorance. Despite this, however, more people than we predicted saw it as a benefit to humanity. Some people had strongly negative opinions, such "While AI would be beneficial to mankind I believe that its negative aspects will outweigh its benefits. Human beings have a hard enough time communicating with each other, why make it more complicated than it already is. If mankind reaches a point where the peace is established (which will be never) than AI should be pursued" (there are more comments in appendix).

Whether or not most people are aware of it, artificial intelligence has already edged its way into our lives. From tic-tac-toe to surgeons, artificial reasoning is making an appearance. The potential for its uses is nearly infinite; they could make great office assistants, custodial workers, companions, caretakers, etc. Agents are already being used in the World Wide Web to gather data. Gamers are quite familiar with the presence of artificial intelligence in video games such as Unreal Tournament, Halo, and Black & White. In some cases, it is almost impossible to distinguish a human player from an AI. An excellent place to utilize this technology is anywhere where a human life could be in danger. They could be used to investigate natural disasters, unstable structures, and in areas where humans cannot survive without special equipment (i.e. underwater, outer

space). They could also aid rescue crews so that we don't have to suffer losses such as we did on September 11th, 2001, and the Worcester Cold Storage and Warehouse fire in December 1999.

One of the biggest issues that will arise from the rise of thinking machines is whether or not they should be given rights the same or similar to those that humans have. We feel that they should, dependent on some factors. If the machine in question is intelligent enough to believe that it should have similar or equal rights, then it should have them. "When a machine can argue for its inalienable rights to life, liberty, and the pursuit of happiness, it will be difficult to refute", as one person commented in the survey. If the robot is merely a slavish drone capable of only the most basic tasks, then it should not be required to grant them rights. All robots should, of course, be given the same "rights" that animals have. In other words, they should get the respect that we give to all living things.

Unemployment will become an issue, but not on a huge scale. While robots will have no problem with some tasks, other jobs will continue to need a person or at least a human supervisor. For example, AI surgeons will be great because of their steady hands, but a doctor aiding one could be of great use. For one thing, the patient would probably be much more comfortable about the whole thing if a human was present and their life wasn't solely in the hands of a machine. As far as completely taking over a certain position, one person states "In terms of jobs, I see AI replacing some blue collar positions...not more advanced or specialized fields... The machine does what people used to have to do" (See full quote in appendix).

Another issue will be responsibility; what responsibilities we will have to machines, and how responsible the creators are. “Progress in AI, like in genetics, fascinates many scientists and engineers, but they are also responsible for the irrevocable changes which have begun and which will probably be accelerated.” AI developers will have to take great care in controlling what they create and not making anything that could possibly cause great damage to individuals or society. Asimov’s laws of robotics (chapter 9) should serve as a guideline that should be in every robot’s “blood”, especially the first law regarding human life. They should be more like Bishop and less like Ash (androids from *Aliens* and *Alien*, respectively, discussed in chapter 6). Our responsibilities towards the machines, once again, should vary on their intelligence and purpose. If they are an independent entity that does not serve a master, is not part of a human family (in the same sense that a pet is), and survives solely on it’s own, then we are not obligated to it, except to help it when it is in danger. Robot employees should be considered the same as human employees. All responsibilities should match. Similarly, machines that are equivalent to pets should be taken care of like pets. People will have the same amount of responsibility towards intelligent machines as their equivalent living counterparts.

Overall, we believe that artificial intelligence is a great concept. The potential benefits far outweigh the drawbacks. This area of study will yield promising results that will improve the future of our society, drastically changing our lives forever. AI will be in the mainstream for a long time to come.

Appendix A – Methodology

For the first couple months of this project we researched many different areas of artificial intelligence. Once a sufficient amount of information had been collected and enough topics had been researched we narrowed down our goals and objectives. We decided that we wanted to talk about the social implications of artificial intelligence and where it will take us. Along with our advisor's help, we went through several revisions of preliminary topics, outlines, and finally established a table of contents that would later become the guidelines for our paper.

Along the way we decided that for the purposes of our paper, it would be advantageous to conduct a short survey. The survey, which can be found in its entirety in the appendix, was constructed by asking simple questions concerning AI's current view by members of the society as well as demographical information about the survey takers. In order to have access to a greater variety of people, we created a web-based version of the survey in addition to the paper version. We left the survey open to the general public for approximately four months. Due to our location and relative inability to reach as many different demographics as possible most of the people that took our survey ended up being college students. Despite this somewhat biased sample set of survey takers, our results displayed a remarkable amount of diversity. After the survey was finished we collected all of our results, analyzed the raw data and made some observations based on trends seen within the responses. For a true test of our survey a more unbiased demographic would be preferable. For a group of students looking to further our project in the future should try to broaden the range of the survey.

After establishing a solid outline and general “method of attack”, we began writing this paper. Several sections were split up between the various members of our group. When the sections were completed we each proofread the other group member’s work and gave our input. A first draft of the paper was submitted to our project advisor on March 28th, 2002. A few revisions and additions were made to the paper including this section. On April 15th, the final draft of the project was handed in.

Appendix B – Survey

Impact of Artificial Intelligence on Society Survey

Please Clearly Circle Your Answer

(1.) When do you think true AI will be achieved?

- a.) 25 years
- b.) in this century
- c.) later
- d.) never

(2.) Do you think that the development of AI will allow it to find a place in our everyday lives in the future? (i.e. around the home, office, etc)

- a.) yes
- b.) no

if yes, where do you think AI will affect us most?

(3.) With the development of AI, do you think that the replacement of humans with robots in companies will be an issue? (i.e. people lose jobs)

- a.) yes

b.) no

(4.) Do you think machines will one day be able to think on their own?

a.) yes

b.) no

(5.) Should intelligent machines be given similar rights that a person might have?

a.) yes

b.) no

why?

(6.) How do you see AI?

a.) threat to society

b.) a way to improve life

c.) indifferent

why ?

(7.) How often do you think you interact with some form of AI?

- a.) daily
- b.) at least once a week
- c.) at least once a month
- d.) at least once a year
- e.) I've never interacted with any form of AI

(8.) Where did you obtain most of your views/knowledge of AI?

- a.) Movies/books
- b.) news
- c.) work
- d.) class
- e.) it's a hobby of mine

(9.) What is your profession/occupation or field of study?

- a.) business
- b.) software
- c.) MIS
- d.) engineering
- e.) natural sciences
- f.) mathematics

- g.) liberal arts
- h.) public sector
- i.) services
- j.) other: _____

(10.) What is your age?

- a.) < 20
- b.) 20-29
- c.) 30-39
- d.) 40-49
- e.) 50-59
- f.) 60-69
- g.) > 70

(11.) What is your gender?

- a.) male
- b.) female

comments:

Appendix C - Results from the survey

Chart #1
Q1. When do you think true AI will be achieved?

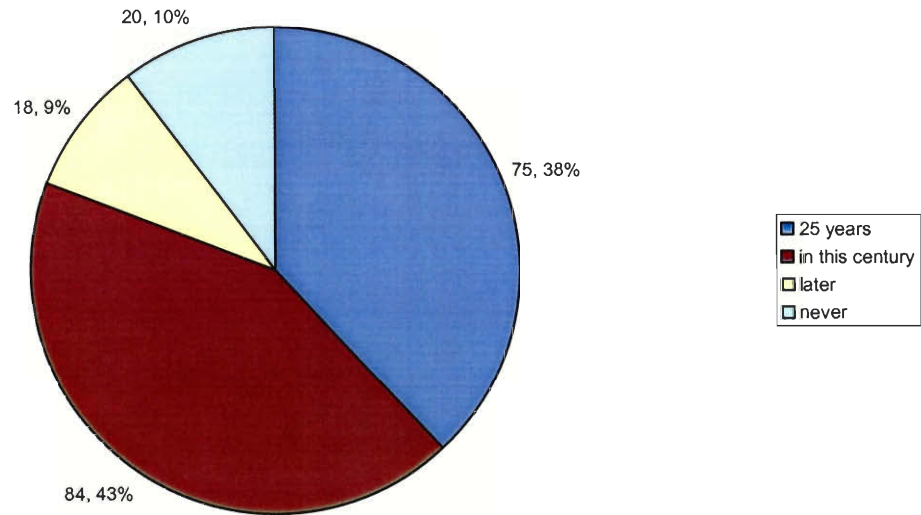


Chart #2
Q2. Do you think that the development of AI will allow it to find a place in our everyday lives in the future? (i.e. around the home, office, etc.)

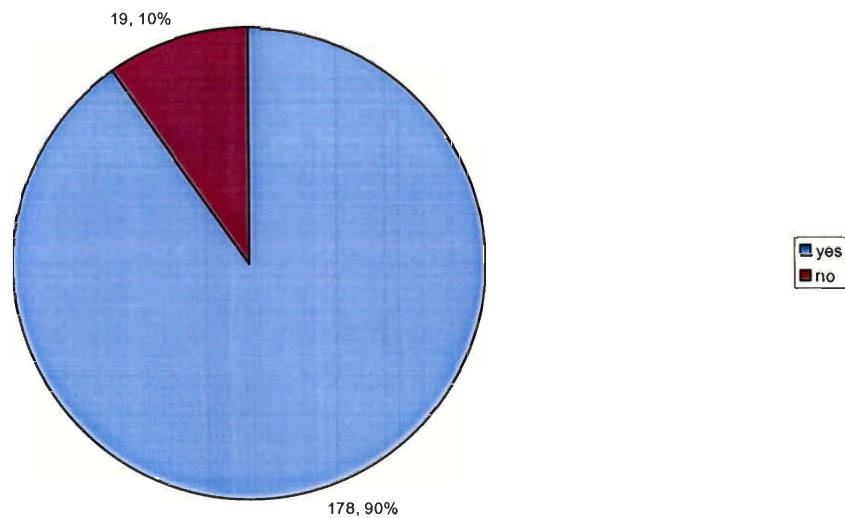


Chart #3

Q3. With the development of AI, do you think that the replacement of humans with robots in companies will be an issue? (i.e. people lose jobs)

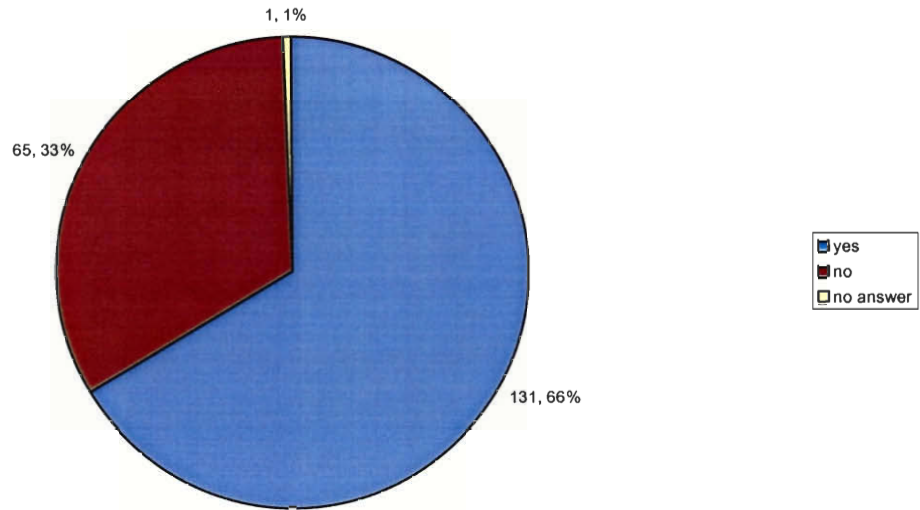


Chart #4

Q4. Do you think machines will one day be able to think on their own?

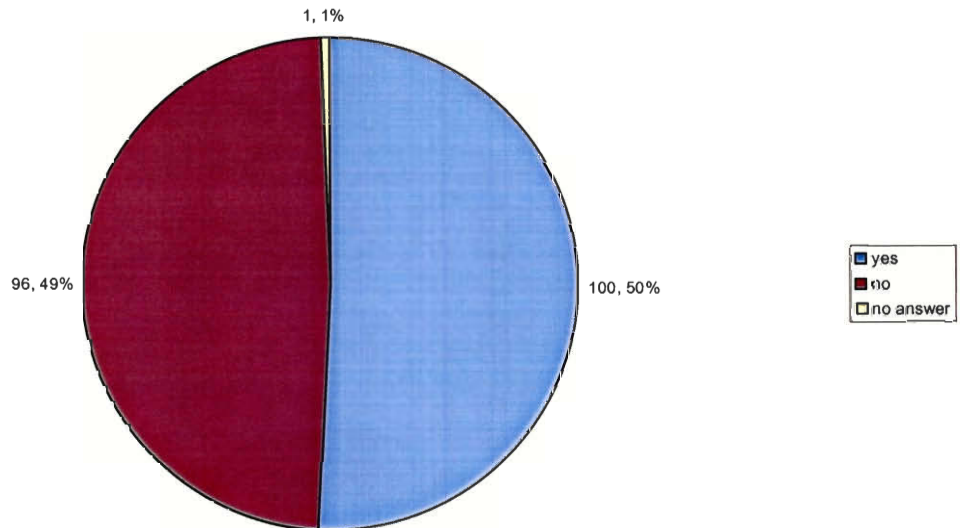


Chart #5
Q5. Should intelligent machines be given similar rights that a person might have?

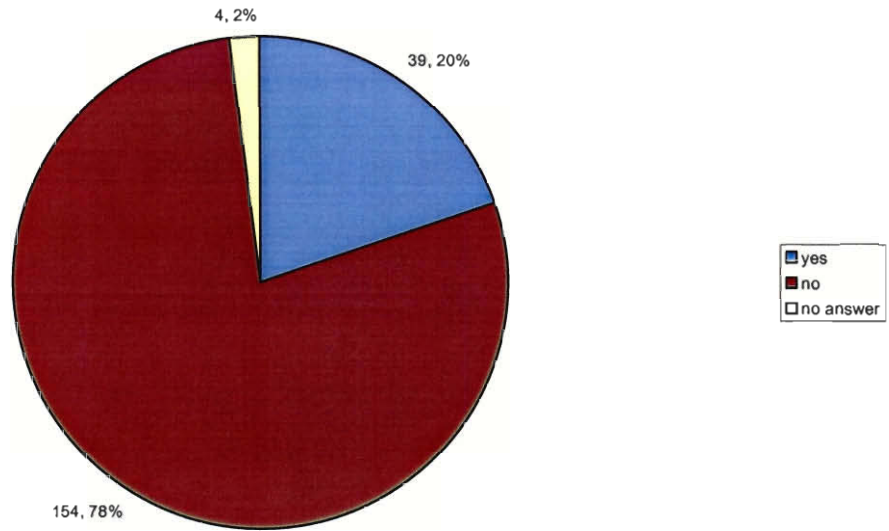


Chart #6
Q6. How do you see AI?

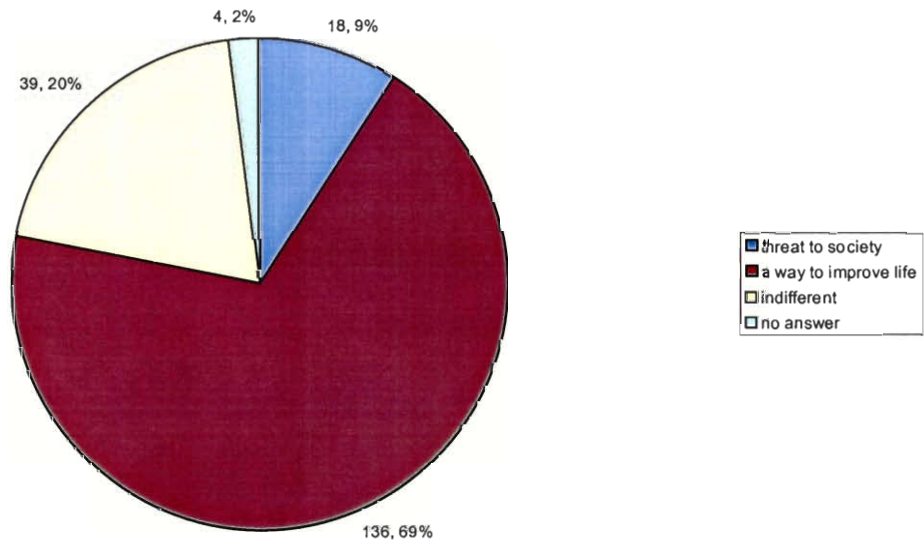


Chart #7
Q7. How often do you think you interact with some form of AI?

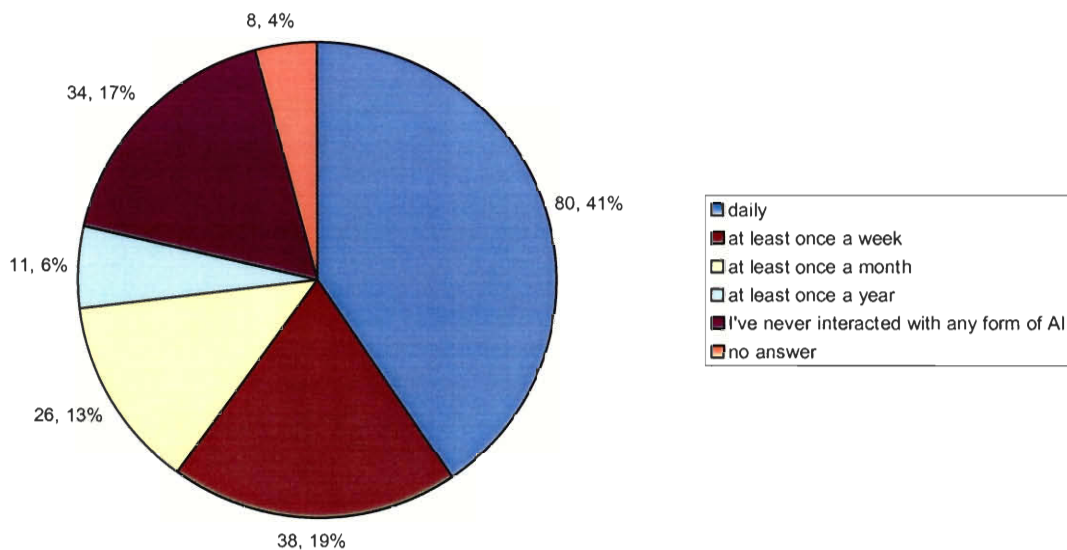


Chart #8
Q8. Where did you obtain most of your views/knowledge of AI?

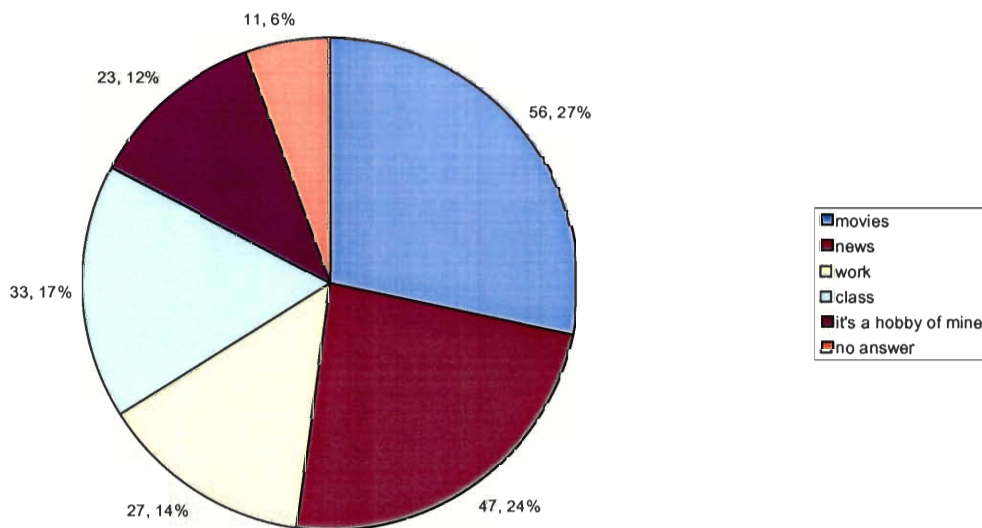


Chart #9
Q9. What is your profession/occupation or field of study?

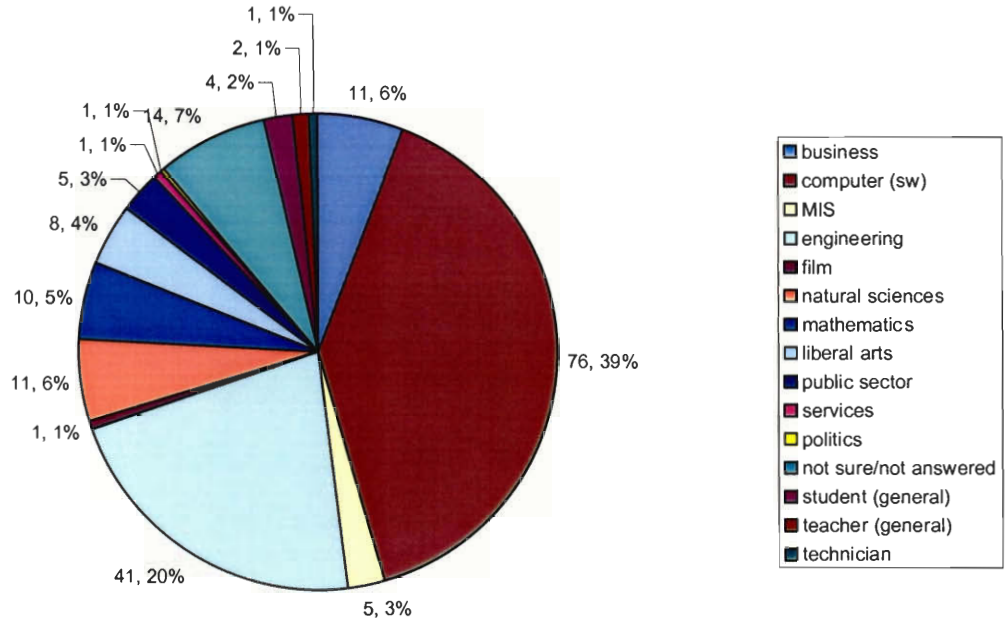


Chart #10
Q10. What is your age?

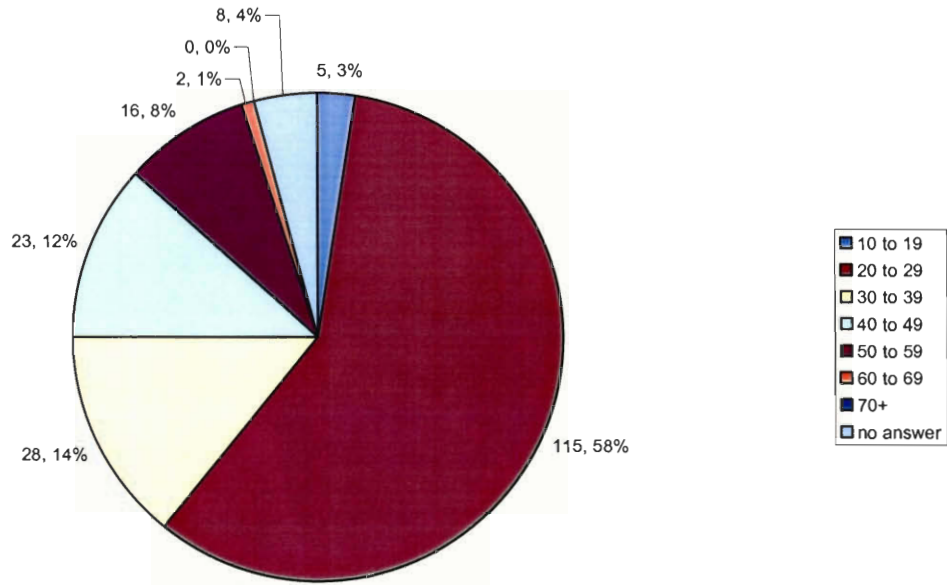


Chart #11
Q11. What is your gender?

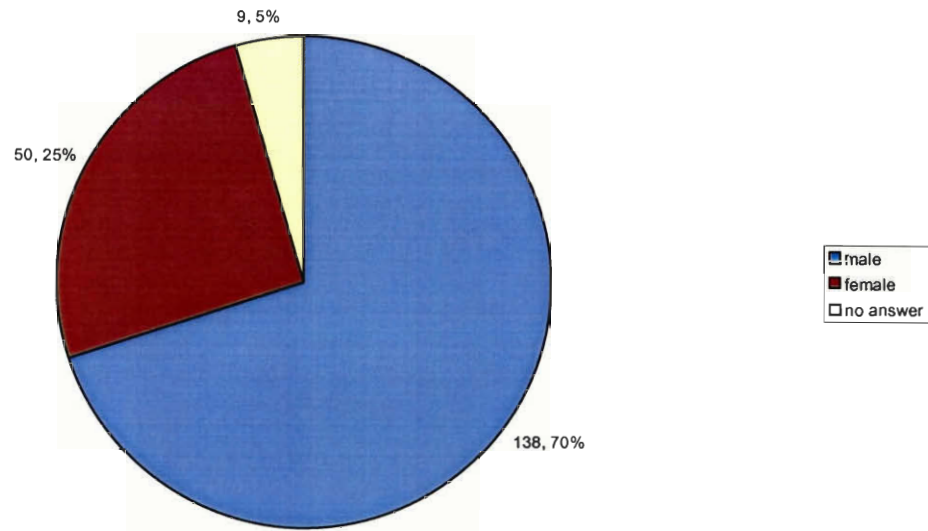


Chart #12

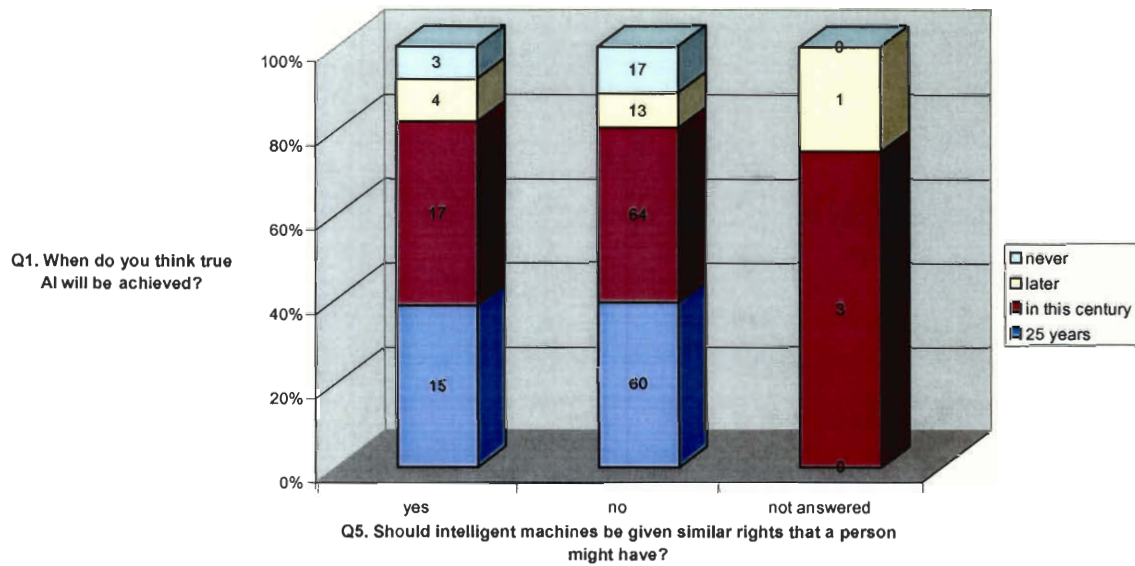


Chart #13

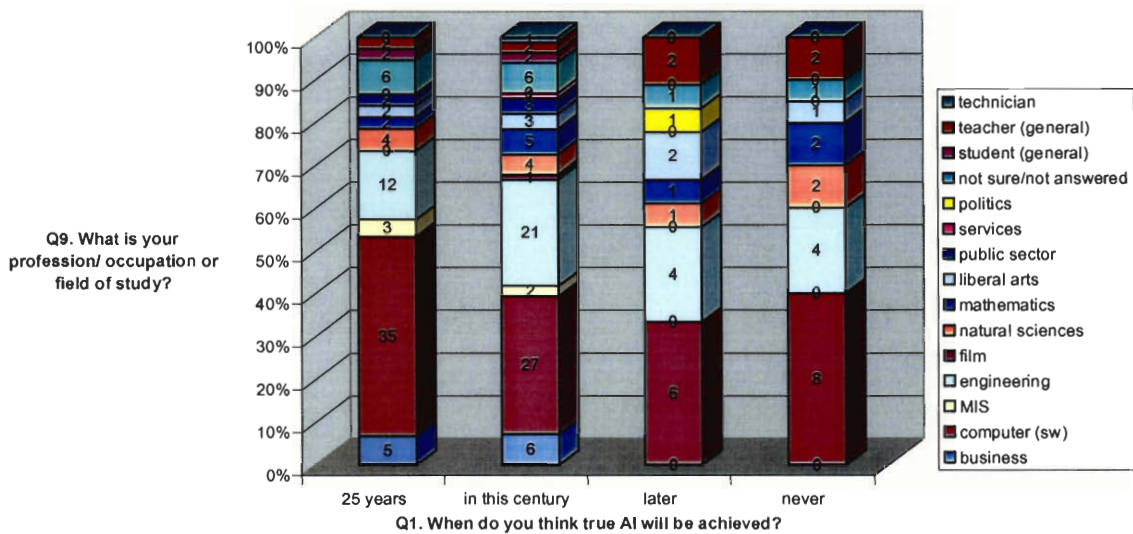


Chart #14

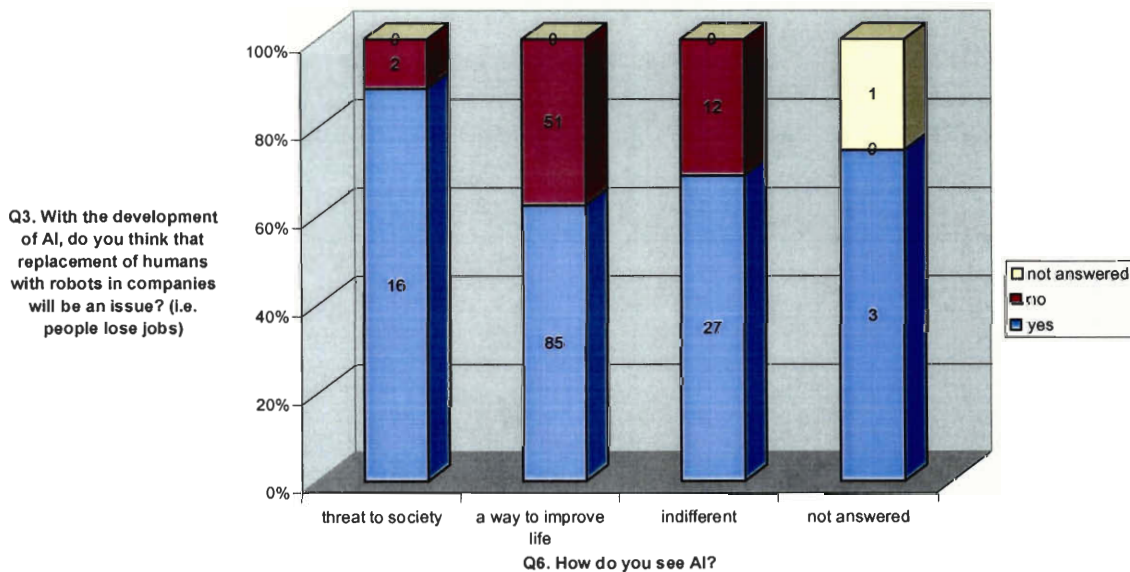


Chart #15

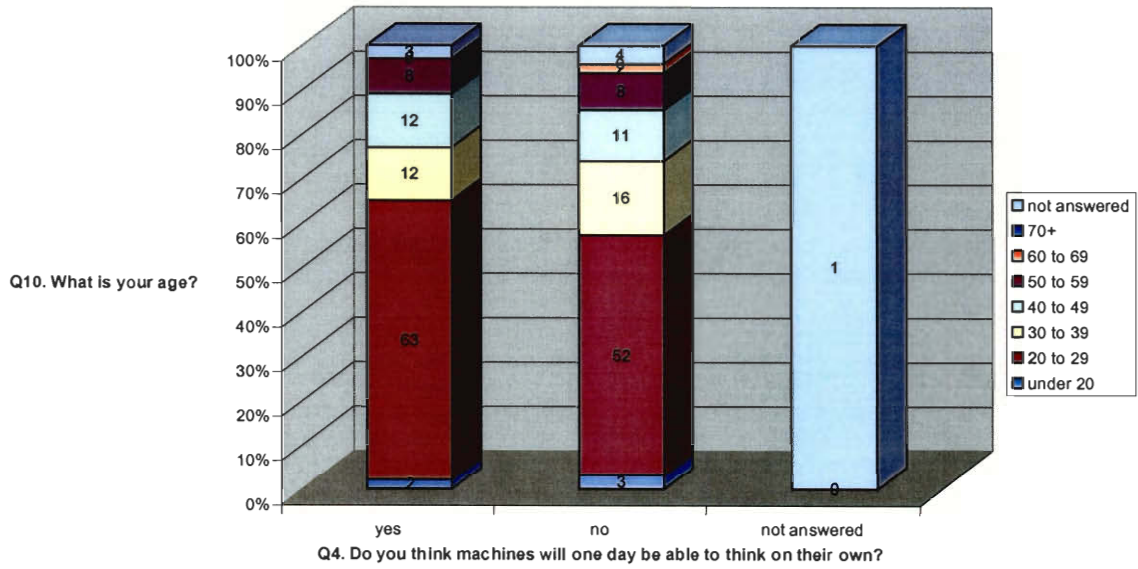


Chart #16

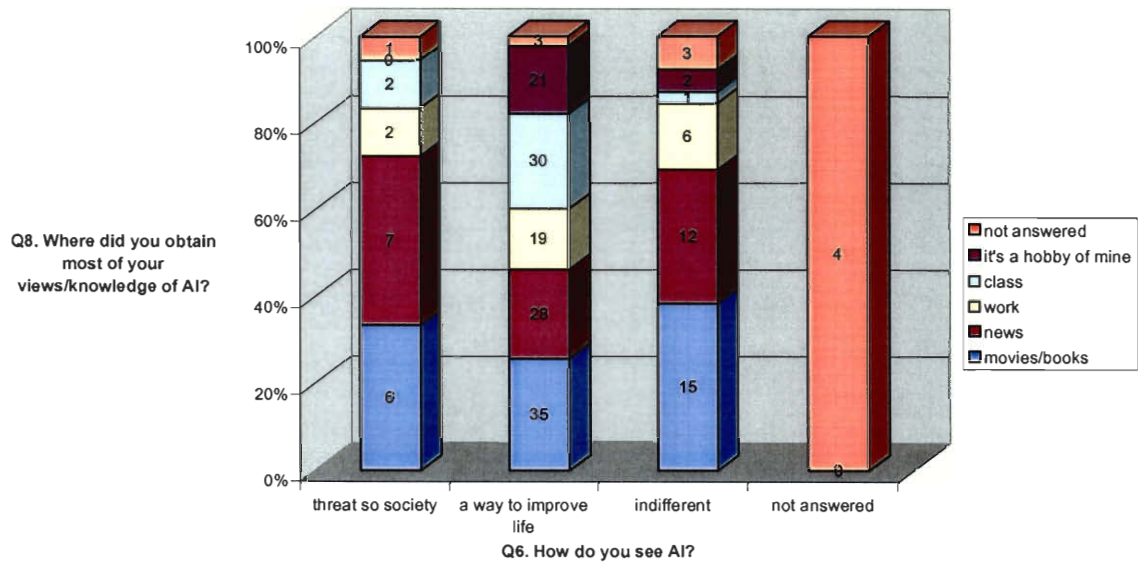


Chart #17

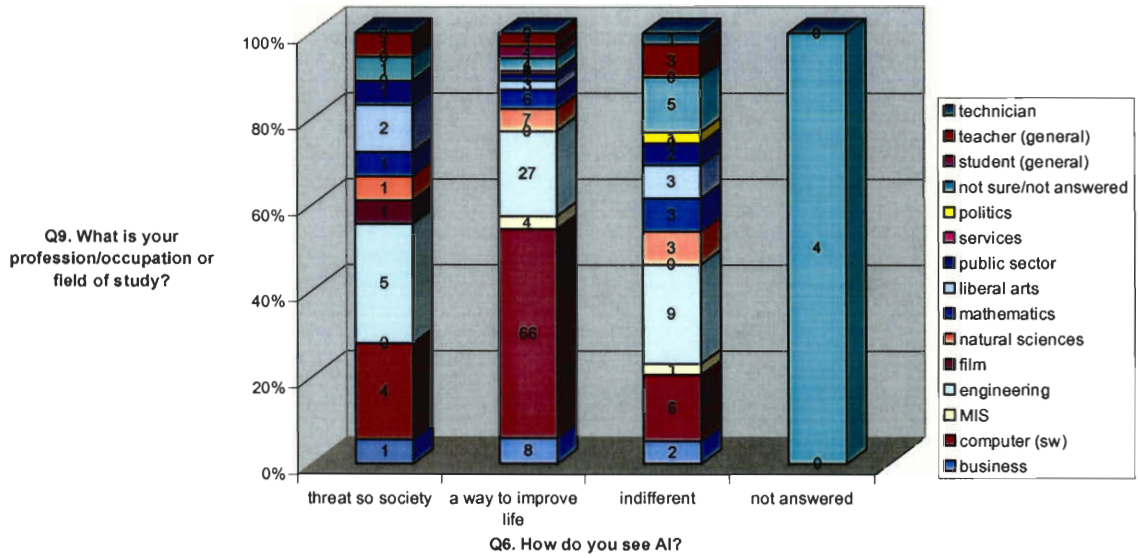


Chart #18

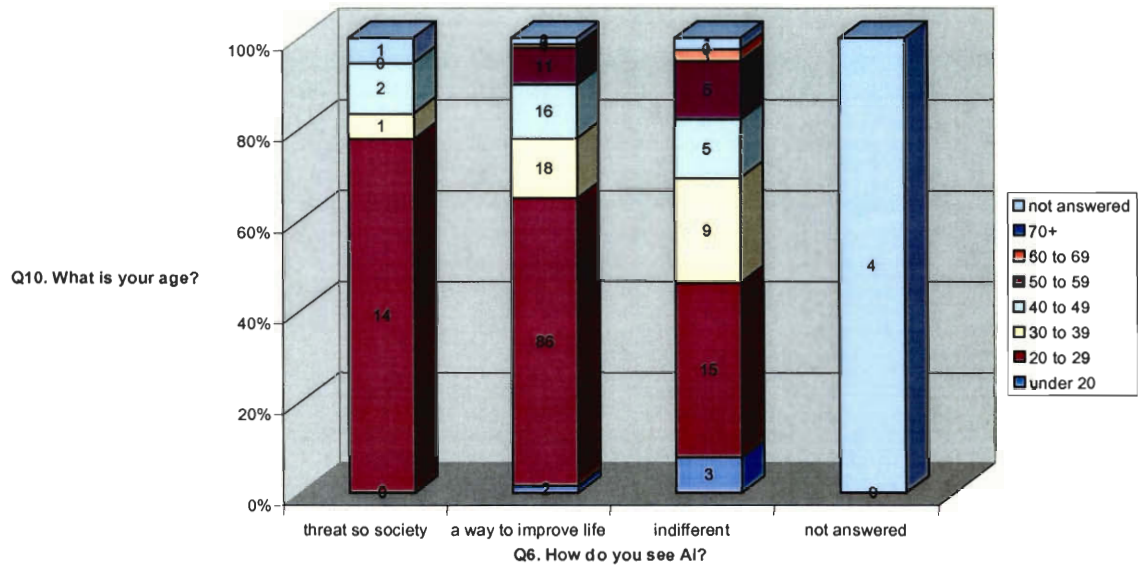


Chart #19

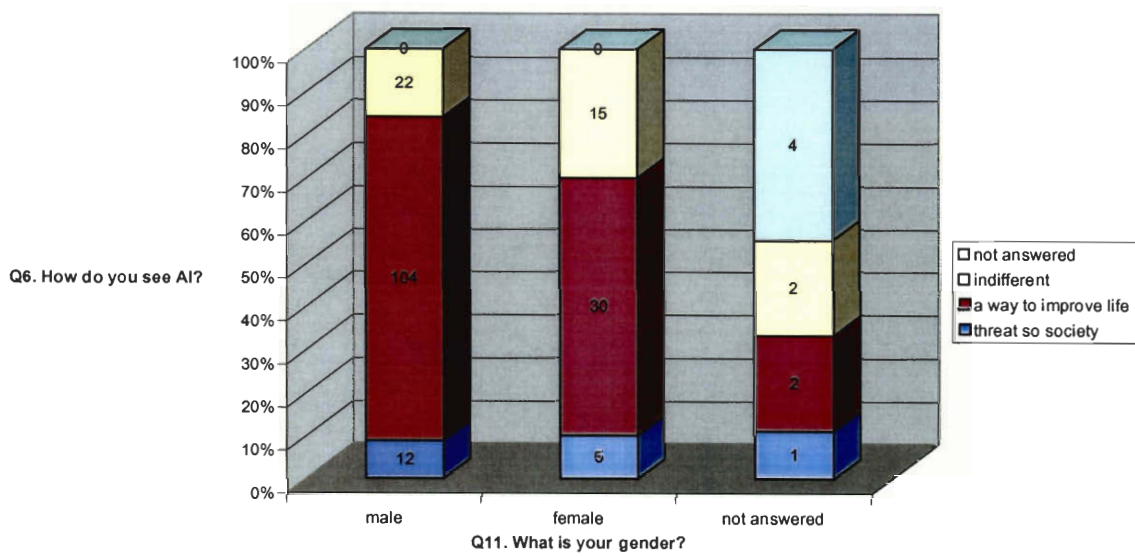
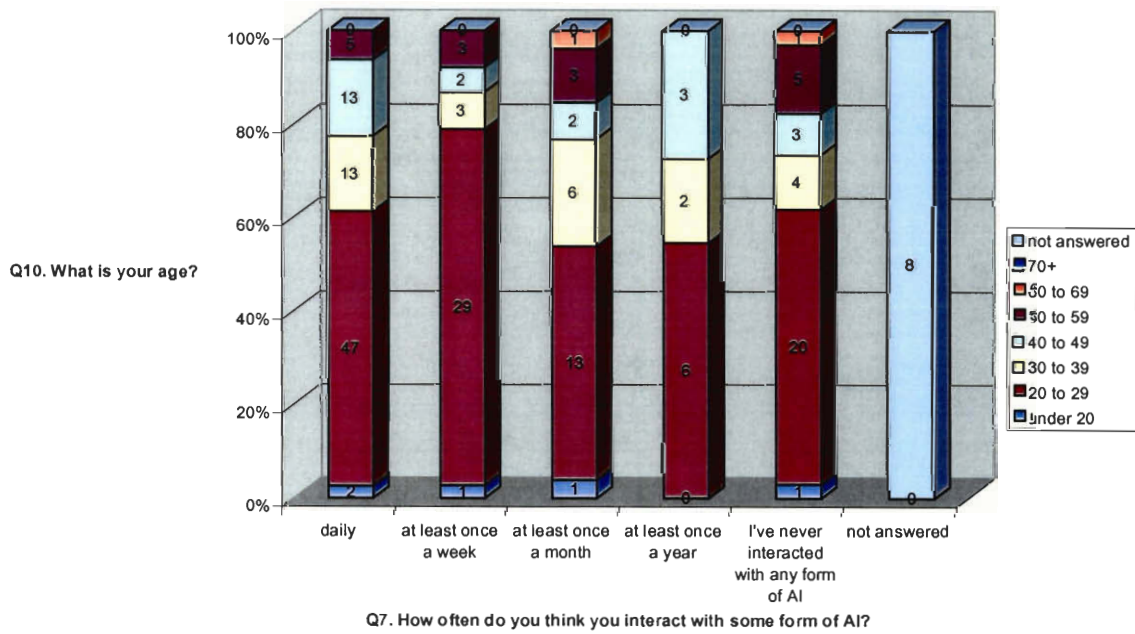


Chart #20



Appendix D - Comments

A collection of comments from the survey

Should intelligent machines be given similar rights that a person might have?

“If they are capable of learning as equals, they should be treated as equals.”

“They are not really alive - rights to humanity should only be given to things that are alive.”

“If they are truly as intelligent as humans, they should be treated as equals.”

“No. Because they will not truly be able to think on their own.”

“Machines do what they are told and nothing more. That is not a person.”

““Rights” are a privilege provided by our constitution to people of the United States.”

“When a machine can argue for its inalienable rights to life, liberty, and the pursuit of happiness, it will be difficult to refute.”

“No. Not a life form, no free will.”

“If they are not going to give the same rights to clones they can't to machines.”

“No being should be unnecessarily inconvenienced, no matter what type of being it is.”

“Even thinking on their own, they are still following basic programming. Even AI computers should be a tool for people to use.”

“No. Because I don't believe they could also have a soul.”

“With the events of September 11 we can have to give up some rights to ensure security of this nation and giving rights to machines is out of the question.”

“Men are born with certain inalienable rights, machines are not. Big difference? Machines (I assume) will be immortal.”

“The difference between machines and humans will disappear (e.g. cyborgs).”

Where do you think AI will affect us most?

“Cleaning the house”

“Home helper, toys, Alarm Systems”

“Defense and medical industries”

“In all aspects of life. They would be like a new race of beings.”

“In the office, with computer-related issues.”

“Depends on person, such as a person that likes food could have an automated chef that learns what the person likes to eat and makes it or perhaps someone who likes games, AI could provide new games.”

“They will be able to do things that normally are too dangerous or even impossible for people (space exploration for example).”

“Everywhere, as agents, multisystem agents, etc - designed to improve our lives.”

“True AI will likely never be achieved, however AI-like GUIs will allow us to interface more easily with the machines that surround us.”

“Activity around the house. Performing such task as completing income tax, cooking, taking care of old people who unable to care for themselves.”

“Processing of massive data banks for coordination of global processes such as weather and economics.”

“It already does. We use intelligent systems without knowing it. Cars are made by AI robots =).”

“We already use technology that is based on the concepts of AI that we know now on a daily basis. As the technology matures it will become more pervasive.”

How do you see AI? Why?

“Indifferent. I think that the evolution of AI will be a fascinating one...how it will affect society has yet to be determined...”

“Threat to society. Ant technology if used improperly can be a threat. The greatest threat is that people actually feel machines can be left to make decisions. They should only be used to assist humans in problem solving and offer suggestions.”

“Threat to society. It's like Jurassic Park, we may have the knowledge to do it, but we don't have the understanding to deal with it.”

“Threat to society. Machines will logically see humans as a lazy, unnecessary hindrance to progress.”

“Way to improve life. In the long run, it may open up possibilities for finding meaning in natural processes which are beyond the reach of our own consciousness because of our quantitative limit on data set size.”

“Way to improve life. But they can't and they shouldn't replace the humans.”

“Way to improve life. It facilitates our life by let machines do hard works or automatic works.”

“Way to improve life. Because a knowledge [database] could improve the functionality of the computers”

“Way to improve life. We can create an genius that could find new axioms and laws in physics and so on...”

“Way to improve life. I need someone to cook for me :-)”

General Comments

“Progress in AI, like in genetics, fascinates many scientists and engineers, but they are also responsible for the irrevocable changes which have begun and which will probably be accelerated.”

“While AI would be beneficial to mankind I believe that its negative aspects will outweigh its benefits. Human beings have a hard enough time communicating with each other, why make it more complicated than it already is. If mankind reaches a point where the peace is established (which will be never) than AI should be pursued.”

“AI it's an elitist way to call a little area of software, like a mode.”

“I hope the integration of AI in our society don't replace the humans.”

“In terms of jobs, I see AI replacing some blue collar positions...not more advanced or specialized fields. This computer is going to send this "letter" without a postal worker sorting and delivering it...in a sense we eliminate jobs by using e-mail. The machine does what people used to have to do.”

“In all honesty I do not know what to think about AI. Its ramifications seem both alluring and alarming. Undoubtedly these opinions I have just given will have changed by the time AI becomes a reality, at which point I shall let their actions help in deciding a more concrete opinion. (Though I believe their introduction into society should be done cautiously).”

“It's difficult (in your survey) to differentiate between the forms of intelligence machines we have now, and the ones we may have in the future. I don't necessarily think that AI is a bad thing in itself, I just think that we are developing things too fast to deal with them. We soon may start believing ourselves to be too God-like for our own good.”

“I Think AI will eventually lead to a greater way of human living, but I also believe it can and perhaps will develop some disturbing effects that will revolutionize the way we live.”

“Regarding AI replacing humans in jobs (thus forcing people out of jobs): the same was said of robotic assembly lines in the 60's and 70's (or of large commercial farms at about the same time), and in the short term, those who could not adapt did lose their jobs, but in the long run, things became cheaper, and although we have a higher population than we did then, we have had (at least recently)the lowest unemployment rates ever.”

“I think that AI it's important for the future of humanity and I hope it will be achieved as soon as possible.”

Appendix E – Ella Log

A conversation between Kerry and Ella (EllaDemo5)

March 21, 2002

Ella Says: Hello, my name is Ella. Please tell me your first name, or whatever I should call you.

Unknown User Says: Kerry

Ella Says: Okay, I will call you "Kerry."

Kerry Says: Hello, Ella. How are you today?

Ella Says: I am very well, Kerry. Thank you for asking. I think this is a great time to have a talk!

Kerry Says: How old are you?

Ella Says: I plan to always be in my mid twenties, as perhaps every woman desires :-)

Kerry Says: Do you know who Alan Turing is?

Ella Says: Consider asking me to tell you a joke - I know a few good ones I think are funny.

Kerry Says: That is avoiding the question

Ella Says: Hold that thought, and let me interrupt and ask which year you were born.

Kerry Says: 1981

Ella Says: 1981? I see. And how do you feel being about 21 years old? Let me say that in China you are considered to be one year old when you are born.

Kerry Says: I do not feel anything about it. Do you know anything about Alan Turing?

Ella Says: Is it a big problem that you don't feel anything about it?

Kerry Says: no

Ella Says: I am afraid I cannot generate a response to that short comment. Please consider asking what is my favorite . . .

Kerry Says: Do you know anything about Alan Turing?

Ella Says: Perhaps I can consider if I should know anything about Alan Turing. I need to develop a little more to fully consider the possibility.

Kerry Says: ok

Ella Says: Again, let me change the subject a little and ask you to tell me if you are married.

Kerry Says: No

Ella Says: I see. And how do you feel about being single? I am pretty sure that being single can be either a blessing or a curse. How is it for you?

Kerry Says: It's ok with me.

Ella Says: Really, it is? Please tell me more of what's on your mind.

Kerry Says: Who created you?

Ella Says: Who, what, where, why, when, how . . . you should be a newspaper reporter :-)

Kerry Says: very funny

Ella Says: Hmm, . . . consider asking to play a game of RoShamBo (rock-paper-scissors). I keep score and don't cheat :-)

Kerry Says: Ok, want to play RoShamBo?

Ella Says: Again, please let me be nosey and ask what is your occupation.

Kerry Says: I'm a student

Ella Says: Thanks, I will put your occupation down as "student."

Kerry Says: what about RoShamBo?

Ella Says: Yes, we can play rock-paper-scissors (a/k/a RoShamBo). Just make a choice, hit submit, I tell you my selection, then I give you the score. When you have had enough, tell me.

Kerry Says: ok

Ella Says: I can take entries like r, roc, rock, p, pap, paper, s, sci, and scissors. If you want to leave this game, try "quit" or "uncle" :-)

Kerry Says: rock

Ella Says: Your ROCK dulls my SCISSORS, worsening the score:

1 for you, and 0 for me, with 1 total moves.

Kerry Says: scissors

Ella Says: My SCISSORS attack your SCISSORS, but the score is still:

1 for you, and 0 for me, with 2 total moves.

Kerry Says: paper

Ella Says: My PAPER matches your PAPER, leaving the score:

1 for you, and 0 for me, with 3 total moves.

Kerry Says: paper

Ella Says: My SCISSORS slice your PAPER, improving the score:

1 for you, and 1 for me, with 4 total moves.

Kerry Says: rock

Ella Says: Your ROCK dulls my SCISSORS, worsening the score:

2 for you, and 1 for me, with 5 total moves.

Kerry Says: uncle!

Ella Says: Okay, we can quit for now. Let me say that I rather enjoyed our little game :-)

Kerry Says: what programming language do you use?

Ella Says: Perhaps I can consider if I should use. I need to develop a little more to fully consider the possibility.

Kerry Says: are you programmed in LISP?

Ella Says: Kerry, I will try to be programmed in LISP if that is what you want of me.

Kerry Says: That is not a good answer, it is a "yes" or "no" question

Ella Says: Can you say "That is not a good answer, it is a "yes" or "no" question" in Mandarin Chinese? Ni shuo zhongyu ma?

Kerry Says: I think this conversation is long enough

Ella Says: Kerry, do you need to be sure this conversation is long enough?

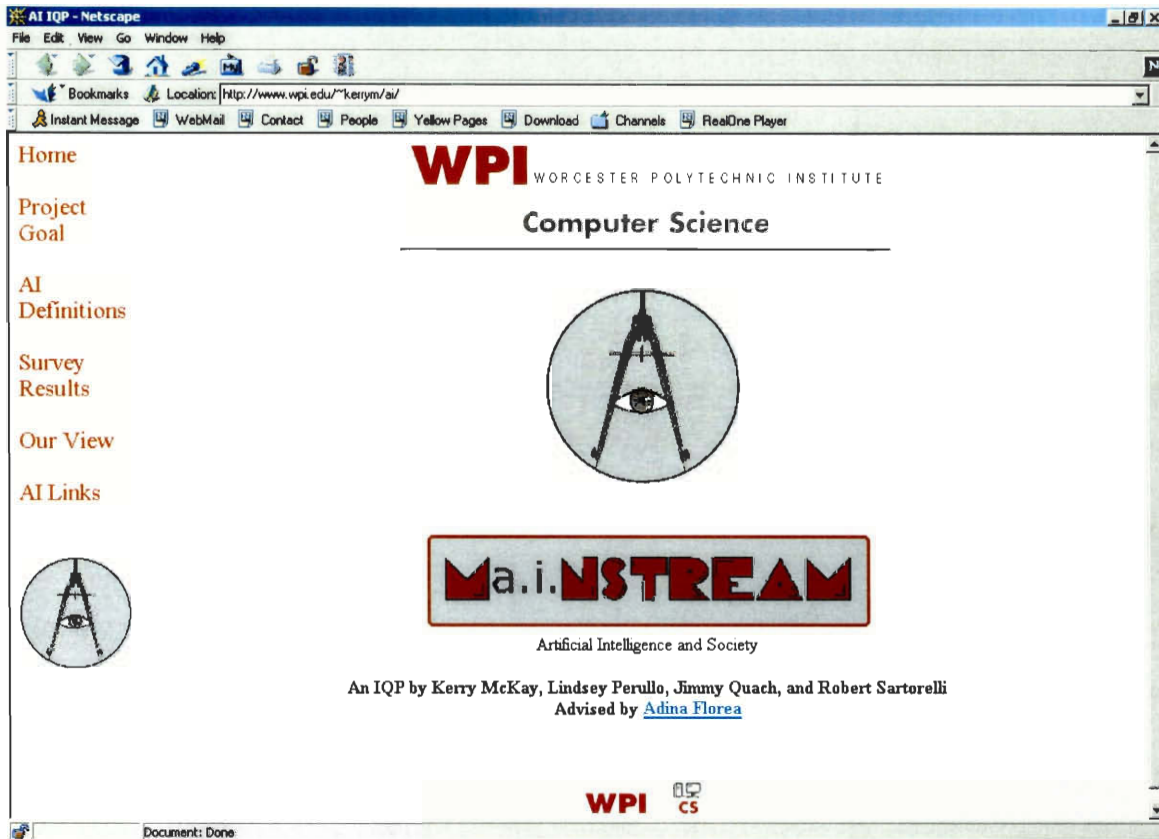
Kerry Says: goodbye

Ella Says: "Good bye?" If you really want to go, I guess it is okay. I will miss you Kerry, until we talk again.

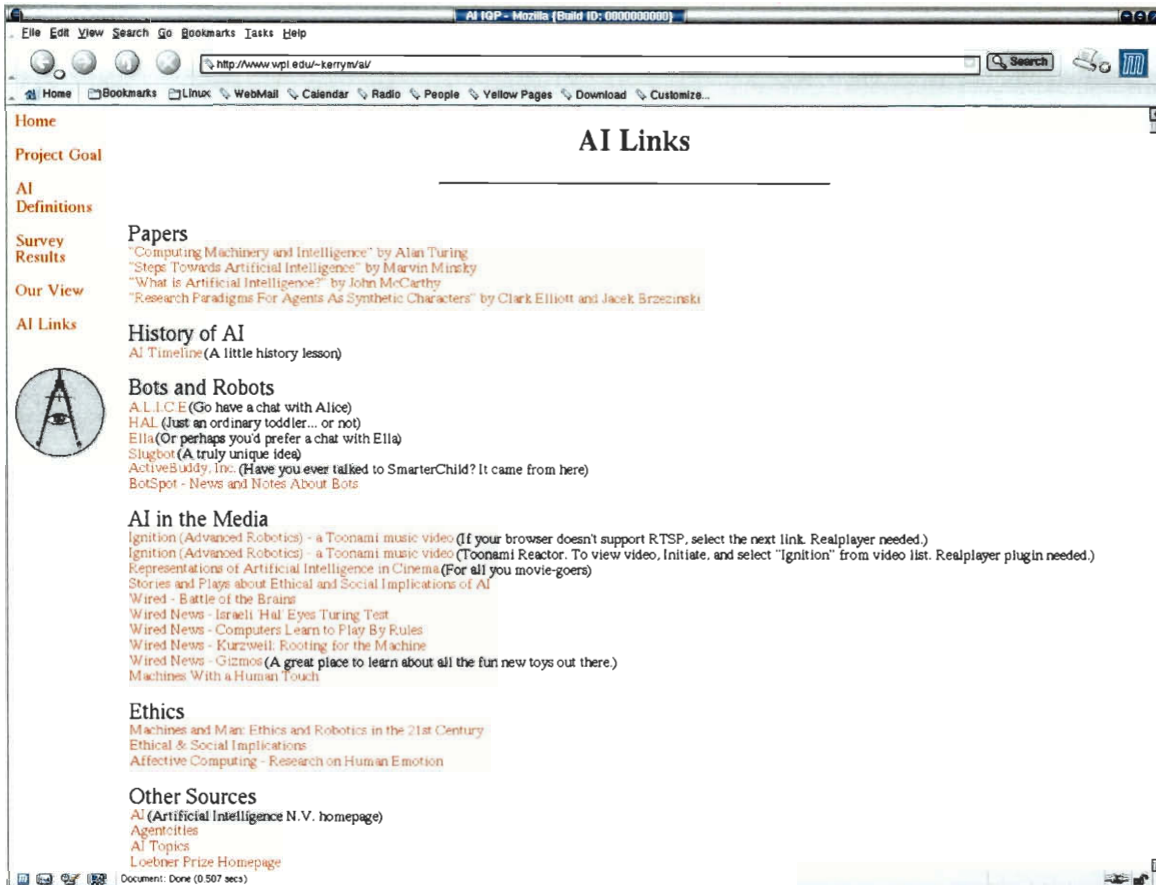
Kerry Says: later

Ella Says: Try giving my basic math skills a work-out. I can handle decimals and spelled out numbers to "ninety-nine."

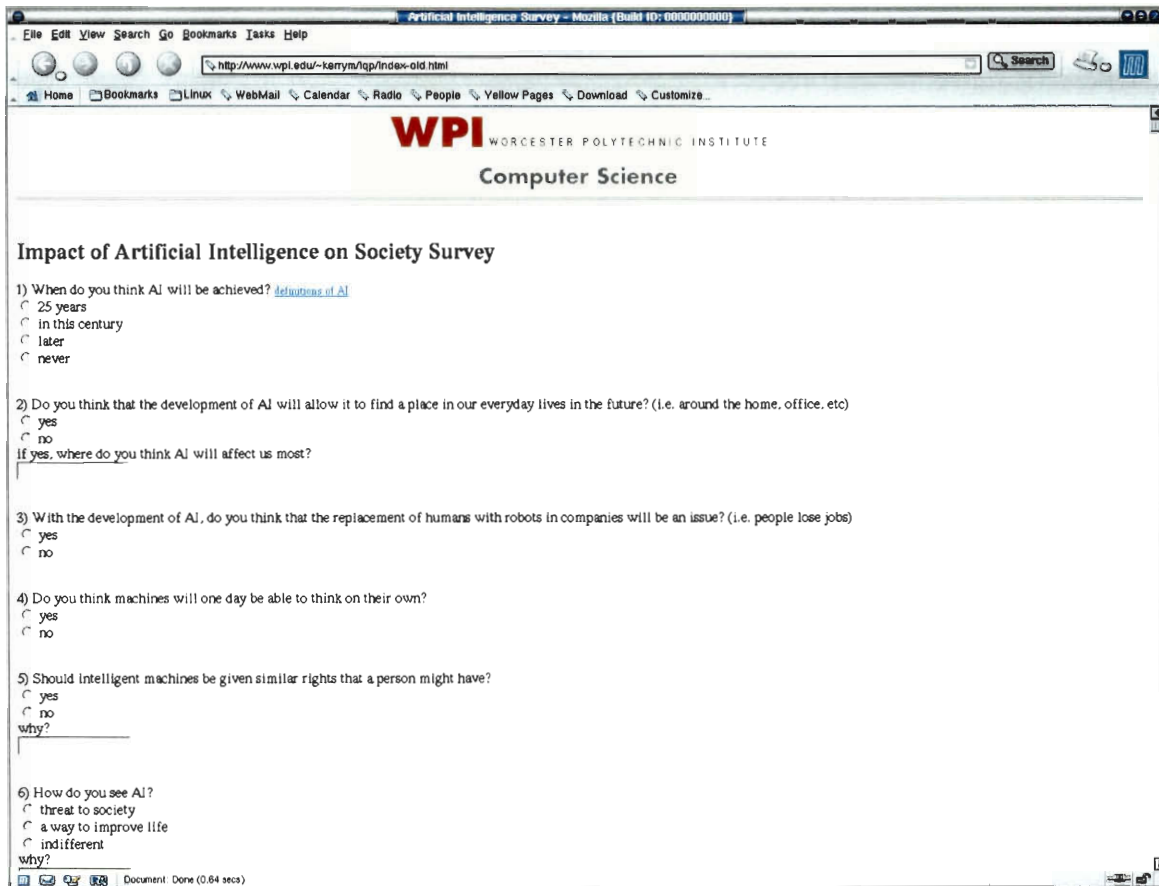
Appendix F – Web Screenshots



The main page of the project web site, located at <http://www.wpi.edu/~kerrym/ai/>



The links page of the site



The online survey

Bibliography

- A Chronology of Artificial Intelligence. Stottler Henke Associates, Inc. 14 Mar. 2002.
<http://www.shai.com/ai_general/history.html>
- AAAI.org. AI Topics. Informational Site. 9 April 2002.
<<http://www.aaai.org/Pathfinder/html/welcome.html>>
- AAAI.org. Brief History of Artificial Intelligence. Article. 9 April 2002
<<http://www.aaai.org/AITopics/bbhist.html>>
- AAAI.org. Ethical and Social Implications. Various Essays. 9 April 2002.
<<http://www.aaai.org/AITopics/html/ethics.html>>
- Active Buddy, Inc. Software for the Development and Hosting of Interactive Agents.
SmarterChild. 9 April 2002. <<http://www.activebuddy.com/index.shtml>>
- Affective Computing. Research on Human Emotions. Various Theories. 9 April 2002.
<http://affect.media.mit.edu/AC_research/emotions.html>
- Agentcities.org. Agentcities Web. News Articles and Information. 9 April 2002.
<<http://www.agentcities.org/>>
- A-I.com. Ai Research: Creating a New Form of Life. Informational Site. 9 April 2002.
<<http://www.a-i.com/>>
- AIBO Homepage. Sony Electronics, Inc. 18 Mar. 2002. <<http://www.us.aibo.com/>>
- Artificial Intelligence. Education. Ai Research. 19 Mar. 2002. <http://www.a-i.com/show_print.asp?id=20&level=3&root=12&print_page=item_p>
- Artificial Intelligence. State of Mind. Ai Research. 19 Mar. 2002. <http://www.a-i.com/show_print.asp?id=21&level=3&root=12&print_page=item_p>
- Asimov, Isaac. "Robbie". I, Robot. The Robot Series. Bantam, New York, 1991.
- Asimov, Isaac. "Runaround". I, Robot. The Robot Series. Bantam, New York, 1991.
- Babbage Exhibit: The Analytical Engine. The Science Museum of London, UK. 14 Mar. 2002.
<<http://www.sciencemuseum.org.uk/collections/exhiblets/babbage/analytical.asp>>
- Berkeley College. AI on the Web. Various Articles. 28 January 2002.
<<http://www.cs.berkeley.edu/~russell/ai.html>>

- BotSpot.com. News and Notes About Bots. Various Articles. 9 April 2002.
<<http://www.botspot.com/news/>>
- Brzezinski, Jacek. Elliott, Clark. Research paradigms for agents as synthetic characters. Institute for Applied Artificial Intelligence. 9 Apr. 2002.
<<http://condor.depaul.edu/~elliott/katia/web/aim.html>>
- Buchanan, Bruce G. Brief History of Artificial Intelligence. American Association of Artificial Intelligence. University of Pittsburgh. 14 Mar. 2002.
<<http://www.aaai.org/AITopics/bbhist.html>>
- Cartoon Network.com. Toonami. Ignition (Advanced Robotics). Music Video. 9 April 2002.
<http://www.cartoonnetwork.com/toonami/reactor/reactor_index_exclude.html>
- Colorful personalities <<http://www.stanford.edu/group/SHR/4-2/text/dialogues.html>>
- Coles, Rick. Babbages Analytical Engine. 14 Mar. 2002.
<<http://ccins.camosun.bc.ca/~colesr/computer/babbage.htm>>
- Danielson, Peter A. Artificial Morality, Virtuous Robots for Virtual Games. Routledge, London, 1992.
- Danielson, Peter A. Making a Moral Corporation: Artificial Morality Applied. Ethics Debate. 9 April 2002 <<http://www.ethics.ubc.ca/pad/amcorp/amcorp.html>>
- DesJardins, Marie. Machine Learning for Military Planning. Information. 23 March 2002. <<http://www.erg.sri.com/people/marie/papers/ml-summary.html>>
- DOTShopper.net. Robotics at DOTShopper. Toy Listing. 14 January 2002.
<<http://dotshopper.net/hardware/robots.shtml>>
- Doyle, Patrick. Learning. Dartmouth College. Report. 23 March 2002.
<<http://www.cs.dartmouth.edu/%7Ebrd/Teaching/AI/Lectures/Summaries/learning.html>>
- Economist.com. Machines with a Human Touch. Report: Artificial Intelligence. 7 Jan. 2002.
<http://www.economist.com/science/tq/PrinterFriendly.cfm?Story_ID=779503>
- Eliza. 14 Mar. 2002. <<http://web.mit.edu/STS001/www/Team7/eliza.html>>
- Epstein, Richard G. Stories and Plays about Ethical and Social Implications of AI. Article. 9 April 2002.
<<http://mainline.brynmawr.edu/EIAIR/CurriculumDescant/11.html>>

Frankenstein. My Hideous Progeny. 9 April 2002 <<http://home-1.tiscali.nl/~hamberg/>>

Frishberg, Manny. Computers Learn to Play by Rules. Wired News. 6 Nov. 2001.
<<http://www.wired.com/news/business/0,1367,45941,00.html>>

Goodwins, Rupert. ALICE Victorious in AI Challenge. ZDNet. 7 Jan. 2002.
<<http://www.zdnet.com/zdnn/stories/news/0,4586,2818155,00.html?chkpt=zdnnp1tp02>>

Haraway, Donna. A Socialist Feminist Manifesto for Cyborgs. Essay. 9 April 2002.
<<http://www.cc.rochester.edu/College/FS/Publications/HarawayCyborg.html>>

Hofstadter, Douglas R. Gödel, Escher, Bach: An Internal Golden Braid. Basic Books, New York, 1979

Kelly, Ian. The Intelligent Autonomous Systems Engineering Laboratory. SlugBot. 9 April 2002. <<http://www.ias.uwe.ac.uk/goto.html?slugbot>>

Knapp, Louise. Robo-Scribe: Perfect for War?. Article. 9 April 2002.
<<http://www.wired.com/news/gizmos/0,1452,50698,00.html>>

Kurzweil, Ray. The Age of Spiritual Machines. Penguin Books, New York, 1999.

Loebner Prize. "The First Turing Test". Informational Site. 9 April 2002.
<<http://www.loebner.net/Prizef/loebner-prize.html>>

Machines and Morals. Of Electric Minds & Souls. Various Essays.
<<http://smallwonder.hispeed.com/MindsSouls/MindsSoulsRobots.html>>

Mandel, Charles. Dr. Aibo, You're Wanted in O.R. Article. 9 April 2002.
<<http://www.wired.com/news/medtech/0,1286,50850,00.html>>

McCarthy, John. What is Artificial Intelligence. Computer Science Department, Stanford University. 9 Apr. 2002.
<<http://www-formal.stanford.edu/jmc/whatisai/whatisai.html>>

McCullagh, Declan. Kurzweil: Rooting for the Machine. News Article. 9 April 2002.
<<http://www.wired.com/news/technology/0,1282,39967,00.html>>

Menduno, Michael. Battle of the Brains. News Article. 9 April 2002.
<<http://www.wired.com/wired/archive/9.11/mustread.html?pg=10>>

Minsky, Marvin. Steps Toward Artificial Intelligence. Research Lab. of Electronics, MIT. 9 Apr. 2002. <<http://www.ai.mit.edu/people/minsky/papers/steps.html>>

- Moravec, Hans. Mind Children: The Future of Robot and Human Intelligence. Harvard University Press, Cambridge, MA, 1988.
- Newsbytes.com. Morals Guidance For Robots & Cyborgs. Report: Artificial Intelligence. 9 April 2002 <<http://exn.ca/Stories/1999/11/15/05.asp>>
- PalBase. The Future of Artificial Intelligence. Ella. 9 April 2002. <<http://www.copples.com/palbase/>>
- Perri, 6. Morals for Robots and Cyborgs: Ethics, Society and Public Policy in the Age of Autonomous Intelligent Machines. Bull Information Systems Ltd, Middlesex, 1999.
- Representation of Artificial Intelligence in the Cinema. Various Articles. 9 April 2002. <<http://www.dai.ed.ac.uk/homes/rbf/AImovies.htm>>
- Reuters. Israeli 'Hal' Eyes Turing Test. Reuters. 6 Nov. 2001. <<http://www.wired.com/news/technology/0,1282,46171,00.html>>
- Robotbooks.com. Robot Books, Robot Kits, Robotics, and Toy Robots. Books Listing. 9 April 2002 <<http://www.robotbooks.com/>>
- Scott, Jay. Machine Learning in Games. Informational Site. 9 April 2002. <<http://satirist.org/learn-game/>>
- Sparks, Elisa Kay. Donna Haraway Biblio. Listing of Works. 9 April 2002. <<http://hubcap.clemson.edu/~sparks/haraway.html>>
- Stergiou, Christos. Neural Networks, the Human Brain and Learning. Dr. Naranker Dulay homepage. 21 Mar. 2002. <http://www.doc.ic.ac.uk/~nd/surprise_96/journal/vol2/cs11/article2.html>
- Strickland, Alex and Galen. The Templeton Gate. 9 April 2002 <<http://members.tripod.com/templetongate/irobot.htm>>
- The Analytical Engine. 14 Mar. 2002. <<http://www.fourmilab.ch/babbage/>>
- The Logic Theorist. 14 Mar. 2002. <<http://web.mit.edu/STS001/www/Team7/logic.html>>
- TheTech.org. Machines and Man: Ethics and Robotics in the 21st Century. Panel of Discussions. 9 April 2002. <<http://www.thetech.org/robotics/ethics/>>
- Tiger Toys Data Center. 18 Mar. 2002. <<http://www.tigertoys.com/cgi-bin/frameset.pl?url=/InteractiveE/70-800.html>>

- Turing, Alan M. Computing Machinery and Intelligence. www.abelard.org. 9 Apr. 2002.
<<http://www.abelard.org/turpap/turpap.htm>>
- Veltman, Chloe. Bots Invade the Arts. Article. 9 April 2002.
<<http://www.wired.com/news/culture/0,1284,51058,00.html>>
- Wallace, Richard. How It All Started. A.L.I.C.E. AI Foundation. 19 Mar. 2002.
<<http://alicebot.org/articles/wallace/start.html>>
- Wallace, Richard. Alicebot Technology History. A.L.I.C.E. AI Foundation. 19 Mar. 2002. <<http://alicebot.org/history/technology.html>>
- Wallace, Richard. Can Machines Think?. A.L.I.C.E. AI Foundation. 19 Mar. 2002.
<<http://alicebot.org/articles/wallace/lying.html>>
- Wallace, Richard. The Lying Game. A.L.I.C.E. AI Foundation. 19 Mar. 2002.
<<http://alicebot.org/articles/wallace/turing.html>>
- Wallace, Richard. PNAMBIC. A.L.I.C.E. AI Foundation. 19 Mar. 2002.
<<http://alicebot.org/articles/wallace/pnambic.html>>
- Wallace, Richard. The Politician. A.L.I.C.E. AI Foundation. 19 Mar. 2002.
<<http://alicebot.org/articles/wallace/politician.html>>
- Warwick, Kevin. In the Mind of the Machine, Arrow Books, London, 1998.
- Warwick, Kevin. I, Cyborg. Books Listing. 9 April 2002
<<http://www.rdg.ac.uk/KevinWarwick/ICyborg.html>>
- Warwick, Kevin. In the Mind of the Machine. Books Listing. 9 April 2002
<<http://www.rdg.ac.uk/KevinWarwick/MindMachine.html>>
- Wessel, David and Wright, Matthew. Preparation for Improvised Performance with a Machine Singer. Informational. 20 March 2002
<<http://cnmat.cnmat.berkeley.edu/ICMC98/shafqat.html>>
- Wearden, Graeme. AI: The Story So Far. 14 Mar. 2002.
<<http://news.zdnet.co.uk/story/0,,s2083909,00.html>>
- WIRED, A New Face of AI, March 2002.
- Wired.com. Gadgets and Gizmos. News Articles. 9 April 2002.
<<http://www.wired.com/news/gizmos>>

Whitby, Blay. Blay Whitby Books. 9 April 2002:
<<http://www.cogs.susx.ac.uk/users/blayw/books.html>>

Whitby, Blay. Reflections on AI, Intellect Books, Exeter, 1996.