TOWARDS GENERALLY INTELLIGENT MACHINES

by

KARAN SHARMA

(Under the Direction of Walter D. Potter)

ABSTRACT

In spite of the impressive successes of many Artificial Intelligence techniques in narrow domains, one of the unrealized goals of AI remains the creation of a generally intelligent machine. It is hypothesized that the *capacity for being creative in a wide variety of environments is a necessary and sufficient condition for general intelligence in machines.* In defense of our proposition, we bring forth three arguments: Law of Requisite Variety Argument, Turing Test Argument, and Complex Systems Argument. Grounded in our hypothesis, we introduce a new test of machine intelligence.

INDEX WORDS: Intelligence, Creativity, Machine Intelligence Test.

TOWARDS GENERALLY INTELLIGENT MACHINES

by

KARAN SHARMA

B.Tech., Kurukshetra University, India, 2005

A Thesis Submitted to the Graduate Faculty of The University of Georgia in Partial Fulfillment of the Requirements for the Degree

MASTER OF SCIENCE

ATHENS, GEORGIA

2008

TOWARDS GENERALLY INTELLIGENT MACHINES

by

KARAN SHARMA

Major Professor: Walter D. Potter

Committee: Khaled Rasheed
Nash Unsworth

Electronic Version Approved:

Maureen Grasso
Dean of the Graduate School
The University of Georgia
December, 2008

ACKNOWLEDGEMENTS

Thanks to Dr. Walter D. Potter for letting me pursue my personal research interests and for providing me with independence to discover my own niches. His comments and advice were extremely useful in shaping this thesis.

Thanks to Dr. Khaled Rasheed and Dr. Nash Unsworth for being the members of my committee and for providing invaluable support whenever necessary.

Finally, thanks to my friends Rashawn Knapp, Benjamin Barber, Will Landecker and Xia Qu for proofreading my thesis in its later stages.

TABLE OF CONTENTS

CHAPTER 1

INTRODUCTION

The single-most important attribute that our species is proud of is our general intelligence. No other species has evolved this remarkable trait. Once evolved, no other living entity could challenge us in the evolutionary game. What exactly is this quality, general intelligence? Psychologists call this factor g; however, there is no precise definition of g. It involves learning, reasoning, problem-solving, creating and so on. The g is a general factor of mental proficiency, such that those who do well on one cognitive task are likely to do well on other cognitive tasks. The g is known to exist in all human-beings; what is different is the matter of degrees not of type. Can we have something analogous to g in machines? This thesis is a theoretical attempt to advance the understanding of general intelligence in machines.

In the past, many efficient AI techniques had successes in narrow domains, but the age-old question remains: What does it take to create the generally intelligent machine? It is hypothesized that the *capacity for being creative in a wide variety of environments is a necessary and sufficient condition for general intelligence in machines.* To this end, we propose a corresponding definition:

Intelligence is the capacity for being creative in a wide variety of environments.

This definition is inspired by the informal definition of intelligence introduced by Legg and Hutter (2007) that says that an agent can be classified as intelligent if it has the capacity to achieve goals in a wide range of environments. We believe that for an agent to maximize its rewards in a wide variety of environments, it must be capable of displaying creative behavior in a wide range of environments. The creativity, as defined here, is to be understood as routine creativity rather than exceptional creativity.

Grounded in this definition of intelligence, we suggest a new test of machine intelligence. A generally intelligent machine that has the capacity to be creative in n domains where n is a large number, must be considered intelligent. This test has both objective and subjective character; however, an objective measurement of creativity is not hard, as has been demonstrated by Ritchie (2001).

The thesis is organized as follows. In chapter 2, we briefly introduce some basic preliminary concepts. In chapter 3, we review intelligence and creativity in the context of both humans and machines. In chapter 4, we argue that creativity can exist without intelligence in computers, but not vice-versa. In chapter 5, we support a new definition and hypothesis with corresponding arguments. In chapter 6, a new test of machine intelligence is proposed. In chapter 7, it is argued that to be creative in a wide array of domains, the way a system manages its knowledge must be a complex adaptive system.

CHAPTER 2

PRELIMINARIES

In this chapter, we outline a brief overview of the concepts discussed in various sections in this thesis.

2.1. Generally Intelligent Machines

Recently, interest has revived in the early goal of AI that aimed to create general intelligence in machines, that is, intelligence in its broader form rather than the traditional narrow approaches. It was in the early 1980s that most researchers digressed to narrow domains of AI (like expert systems, neural networks, logic, genetic algorithms), because it was easier to succeed in those projects. But now, many people are slowly returning to the original goal of AI. The paradigm is challenging and requires massive effort on the part of many researchers. Some researchers have coined the term Artificial General Intelligence (AGI), the paradigm introduced to create in machines, the counterpart of general intelligence in humans. There are many challenges to be resolved – theoretical, mathematical, computational, ethical, philosophical and technological before we get any closer to the authentic general intelligence.

John Searle coined the terms “strong AI” and “weak AI.” Strong AI assumes that the mind is just a computer and that an appropriately programmed computer can possess human-level general intelligence. This hypothesis asserts that the phenomena of the mind can be reduced to symbols. On the contrary, weak AI is a more conservative approach that assumes that computers can be appropriately used in the study of the mind. Searle considered weak AI as a plausible approach and dismisses strong AI on the grounds that mere symbol manipulation cannot account for the mind.

Two other prominent and rival hypotheses are the Physical Symbol Systems Hypothesis (PSSH) and the Physical Grounding Hypothesis (PGH). The PSSH argues that a physical symbol system is a necessary and sufficient condition for intelligence. The proposal states that any computer can be made intelligent by an appropriate symbolic system and that the human brain is also a symbol system. Searle opposed the hypothesis arguing that mere symbol manipulation cannot explain the phenomena of understanding, mind, consciousness and intentionality. Another rival of the PSSH is Rodney Brooks. He introduced the Physical Grounding Hypothesis, which states that for a system to exhibit intelligent behavior, its internal representations must be grounded in reality. The system should form its representations by interacting with its environment through sensors. However, there is no shortage of critics of the grounding hypothesis. One of the major criticisms (Nilsson 2007) has been that at present, the robots based on this hypothesis are just too simple. Brooks still needs to demonstrate the validity of his hypothesis on complex and hard problems.

In the last three decades, many attempts have been made to create general intelligence in machines. One of the earliest attempts, consistent with the symbolic systems hypothesis, was project CYC (Lenat & Guha 1990) that aimed to create common-sense by explicitly declaring knowledge to the computer. CYC contains immense real world knowledge, represented in predicate-calculus and it draws inferences from that knowledge. As of now, the project is considered a failure by most AI researchers.

Another significant attempt is Novamante (Goertzel & Pennachin 2007). The project has its theoretical basis in Goertzel's psynet model. His psynet model claims that the mind is a dynamical system consisting of patterns. Patterns may combine or mutate with each other to yield new patterns. Also, a pattern may associate with other patterns. Some patterns may become habitual if they receive positive reinforcement while others may get destroyed. The Novamante engine, grounded in psynet, is highly complex and relies on evolutionary programming, probabilistic reasoning and logical inferences, besides other techniques.

2.2. Law of Requisite Variety (Ashby's Law)

Ashby's Law (Ashby 1958), well-known in the cybernetics community, states that the variety in a controller must at least be equal to the variety in the controlled. An environment generates a certain set of disturbances that hinder the controller from achieving its goal state. In order to mitigate the effects of these disturbances, a controller must generate counteractions. This law is not just confined to cybernetics or control systems, but has tremendous implications for any system that needs to achieve its goals in a complex environment. For example – imagine an ideal (hypothetical) antivirus system. If this system is to always achieve a successful outcome, it must have as much variety in its defenses against computer viruses as the number of unique viruses.

2.3. Reinforcement Learning

Reinforcement Learning is a machine learning technique in which an agent interacts with its environment. An agent perceives its environment and generates a corresponding action in the environment. If that action aids in goal-achievement, the agent receives positive rewards from the environment; whereas if the action leads to a negative outcome, the agent receives negative rewards. Based on these rewards, the agent learns those actions that aid it in the goal-achievement, while making a tradeoff between exploration and exploitation. Reinforcement learning is a popular technique exploited by many researchers for many complicated applications like robotics, autonomous vehicle control and navigation, and game-playing.

Many ideas in this thesis have been discussed in the reinforcement learning type of framework. This framework, in the context of this thesis, just implies that the agent is trying to maximize its rewards in a particular environment. Such an agent-environment framework was rigorously and formally used by Legg and Hutter (2007).

2.4. Complex Adaptive Systems

Complex Adaptive Systems (CAS) are systems that involve interactions among various components; and, as a result of these interactions, global behavior emerges in a system. The emergent behavior has a disconnect from the local behavior. For example – intelligence cannot be explained in terms of neurons. Examples of complex systems include the human brain, ant colonies, weather and the stock market.

Complex adaptive systems (CAS) are known for various characteristics; among them widely discussed are emergence, self-organization and non-linearity. These systems can mostly not be understood using reductionist techniques, hence, making it harder to comprehend sometimes. Another important property of CAS is non-linearity. That is, it is not possible to sum up the behavior of components to get the resultant, or in other words, the outcome cannot be obtained by linear summation of components. The behavior of a CAS agent is still not understood and many researchers are directing their efforts to it. When we talk of an agent in this thesis, we assume that an agent is a CAS agent. This is obvious because we know that human cognition is a CAS, so any generally intelligent system is likely to be a CAS.

CHAPTER 3

INTELLIGENCE AND CREATIVITY

This chapter reviews the theories and tests of intelligence and creativity. Intelligence and creativity are reviewed in the context of both humans and machines. However, this review is concise, not exhaustive, as it is not possible to do justice to concepts as complex as intelligence and creativity in a single chapter. Yet, every effort has been made to introduce some of the most important ideas.

3.1. Intelligence

Intelligence as a concept has intrigued researchers for more than a century. After decades of research spanning various disciplines such as psychology, philosophy, neuroscience, computer science and AI, the question "What exactly is intelligence" has not been precisely answered. Yet, we all have an intuitive understanding of intelligence. Almost everyone knows intelligence has something to do with learning, problem-solving, creating, reasoning, information-processing and so on. Over the years, many definitions, tests, and theories of human and machine intelligence have been proposed. Even after decades of experimentation and subjugation to theoretical scrutiny, most of these concepts still remain debatable and highly controversial.

Here are three well known definitions:

"Intelligence is the aggregate or global capacity of the individual to act purposefully, to think rationally and to deal effectively with his environment." Wechsler (1944)

"Intelligence is the computational part of the ability to achieve goals in the world. Varying kinds and degrees of intelligence occur in people, many animals and some machines." McCarthy (2004)

"*intelligence is: 1) the ability to achieve one's goals in life, given one's sociocultural context; 2) by capitalizing on strengths and correcting or compensating for weaknesses; 3) in order to adapt to, shape, and select environments; and, 4) through a combination of analytical, creative, and practical abilities*." Sternberg (2005)

One of the most popular theories of intelligence is the general intelligence hypothesis, popularly known as g-factor, introduced by Spearman in the early decades of the twentieth century. The g hypothesis asserts that intelligence is general, and that the same general component underlies all mental capacities. He arrived at this hypothesis after analyzing vast amounts of data on various cognitive tests. It was found that the performance of most people generally remains the same across a variety of different tests (verbal, reasoning, mathematical). A positive correlation was found among various abilities and, hence, performance on one test was predictive of performance on another test. Since then, many psychologists have brought evidence for the g-factor to the fore.

On the other extreme is another faction of psychologists that believes intelligence is composed of many specific cognitive capacities (Guilford 1967, Gardner 1993). Gardner (1993) proposed the theory of multiple intelligences. According to him, there are multiple types of intelligences, viz., verbal, logical-mathematical, musical, bodily-kinesthetic, interpersonal, intrapersonal, naturalist and spatial. The multiple intelligences hypothesis argues that the possession of one kind of intelligence is not a guarantee of the possession of another kind of intelligence. A person high in verbal ability may not excel in mathematical ability.

Midway between the above schools of thought is Robert Sternberg's theory of successful intelligence (Sternberg 2005). According to this theory, an individual is intelligent if he manages to succeed in his own socio-cultural context. The theory proposes that individuals possess three kinds of intelligence - analytical, creative and practical. Since very few possess all these abilities, successfully intelligent people are the ones who have learned to best optimize their abilities. They are the ones who cash in on their strengths and correct or compensate for their weaknesses in order to achieve the desired outcomes. Successfully intelligent people are adept at adapting to an environment, selecting the desired environment and shaping the environment in order to achieve the desired outcomes. After closely scrutinizing Sternberg's claims, Gottfredson (2003a) criticized him for not presenting enough empirical evidence. The theory remains debatable.

However, in spite of the efforts by many researchers to establish theories alternative to g, the g remains the most well-known and strongly established hypothesis. Even among those who have suggested alternative theories, very few completely deny the existence of g. There is convincing evidence that g exists. Beyond doubt, it has been shown that those with high g do better in life than those with low g (Gottfredson 2008). Gottfredson (2003b) argued that high g is advantageous in many of life's outcomes, most prominently in academic performance and job performance. Further, she argued more complex jobs need higher g than simple jobs. Of course, there are domain-specific talents and other personality factors important to success but g remains the single-most significant factor across a wide variety of domains.

The most popular test of human intelligence is the Intelligence Quotient (IQ) test. Most of the time, these tests include verbal, spatial and mathematical questions. The IQ of a person is known to correlate well with academic as well as real life achievement. The predictive validity of these tests is undisputed for the most part by professional psychologists. However, some psychologists have introduced some alternative tests like the tests of Emotional Intelligence which measure the capacity to perceive and manage emotions. Sternberg's tacit knowledge test

(Sternberg & Grigorenko 2001) involves testing the practical aspect of intelligence. In spite of their convincing arguments, none of these tests have been able to replace IQ tests' popularity.

3.2 Tests of Machine intelligence

Many tests of machine intelligence are reviewed in this section. However, for a good survey of machine intelligence tests, the reader is referred to Legg and Hutter (2007).

3.2.1 Turing test

The Turing test is one of the most highly cited and widely debated topics in AI. Since its inception, the Turing test has been discussed in hundreds of papers. The Turing test, proposed by Alan Turing, is a measure of machine intelligence. There is a machine, a human, and a human judge such that the human judge cannot see both the machine and a human. The human judge communicates with both the human and the machine. The purpose of the machine is to make the judge believe that it is human. If the machine succeeds in doing so, and the human judge cannot distinguish between the human and machine, the machine is said to have passed the test.

In his famous Chinese room argument, Searle claims that it is possible that the machine may seem to behave intelligently but what is going on underneath is mere symbol manipulation and the machine does not possess "understanding". It is possible for a machine to pass a Turing test without possessing such deep phenomena as understanding and intentionality.

3.2.2 Total Turing test

The Total Turing test (Harnad 1991) asserts that the ideal test for a machine would be to not just display its human level verbal communication but also human-like robotic behavior. The machine can only be considered intelligent if it not only communicates like humans but it also behaves like them. This test, by being an extension to Turing test, does not remove the fundamental shortcomings of the Turing test (Schweizer 1998).

3.2.3 Truly Total Turing test

The Truly Total Turing Test (Schweizer 1998) demands machines to develop entire sets of cognitive abilities that allows them to attain whatever it takes to replicate the intellectual successes evident in human evolutionary history. In other words, machines, according to this test, must develop basic cognitive capacities that allow the species of robots to create domains like mathematics, chess, language, and so on.

3.2.4 Universal Intelligence Measure

Legg and Hutter (2007) proposed a formal measure of intelligence that computes an agent's capacity to attain rewards in a wide range of environments. If an agent is capable of maximizing rewards in many environments, then it scores high on intelligence. The measure is grounded in rigorous mathematics, assumes a reinforcement learning framework and describes the external working of an agent. However, it is incomputable, hence only theoretically valid.

3.2.5 Psychometric Approaches

Bringsjord and Schimanski (2003) proposed that any machine can be considered intelligent if it is able to perform well on already existing tests of intelligence (like an IQ test). This delivered a new goal for AI – the capacity to pass human psychometric tests. However, this approach emphasizes creating machines that pass tests rather than creating genuine intelligence. What if a machine can only pass an IQ test and is incapable of everything else? Machines have been able to beat humans in chess (which can be considered another version of intelligence test), but they are still not generally intelligent.

3.2.6 Lovelace test

Bringsjord et. el. (2001) introduced the Lovelace test. According to them:

The Agent A, designed by designer D, passes the Lovelace test if

1. A generates output such that this output is not the result of some coincidence but of processes that are repeatable and
2. The designer D could not have predicted the output from A's components (both hardware and software).

Even though it successfully removed some shortcomings of the Turing test, the Lovelace test has some problems. Any machine that can create an output O, even when it is in only one domain, such that the designer could not have predicted the outcome, will be considered successful according to this test. So if a complex evolutionary algorithm is successful in solving some problem, such that it surprises its designers, then it must be designated successful according to the Lovelace test. Hence, the test cannot gauge intelligence in its broader form. But the test has many positive outcomes. The test implicitly emphasizes the notion of creativity in machine intelligence.

3.3 Creativity

The question "what is creativity?" has no simple answer but we all intuitively understand creativity. It is unanimously agreed that it involves the capacity for bringing something new into existence. But not everything that is new can be deemed creative. Almost an infinite number of new entities can be brought into existence. What makes something creative is not just its novelty but also its usefulness. Creativity has been defined as the ability to produce output that is both novel and appropriate (Sternberg & Lubart 1996). Most of the time creativity is associated with music, art, scientific discovery and writing. However, that conception is highly misleading. Creativity not just occurs in every domain but it is the part of our daily lives. All human beings are regularly creative (Pinker 1997, Hofstadter 1985). Coming up with a novel strategy while playing baseball, finding a new way to lure a potential date or decorating a room in an original way are the routine examples of creativity.

Boden (1991) distinguished between P-creativity and H-creativity. P-creativity is known to occur when someone has a new idea with respect to themselves. H-creativity occurs when the idea is of historical importance. For example - when a person generates a new cooking recipe, such that the recipe was not known to him before but the recipe is already known to others, then we can say that he exhibits P-creativity. An example of H-creativity will be when someone discovers a scientific law that was not known before and is historically important. Einstein's discovery of the theory of relativity is an example of H-creativity. Of course, P-creativity is a routine occurrence whereas H-creativity is a rare occurrence.

What exactly is creativity? There are many psychologists and researchers who have wondered over this question for decades. But in the research literature creativity is generally associated with two terms: - novelty and usefulness, there is almost consensus on that. The only problem worrying most researchers is vagueness associated with "new" and "useful." To deal with vagueness of the word "useful", researchers have surrogated the word "useful" with "appropriate." Below, the terms novelty and appropriateness are explained in the context of this paper:

Novelty: Novelty is the quality of being new. Assume an agent is in an environment whose elements can be described by the set $X = \{x_1, x_2, \ldots\ldots, x_n\}$. If a new element y is created by an agent in the environment such that y is dissimilar to all the elements of set X, then y can be said to be novel.

Appropriateness: Appropriateness of an output in an environment can have several different meanings. Fundamentally, appropriateness is that quality of an output produced by an agent that maximizes its rewards (positive reinforcement) from that environment, if we assume the Reinforcement Learning framework. Programs can be trained to measure appropriateness (Ritchie 2001, Ritchie 2007) or creativity can be judged by the designer. It can be the quality of output that lets the agent survive in that particular environment. Also, it can be the comparative value of an output with respect to other similar entities that exist in that particular environment.

3.4 Machine Creativity Tests

Ritchie (2001 and 2007) developed some impressive measures of a programs' creativity. Developing objective measures of creativity that can span across a wide range of domains is hard; nevertheless, he introduced impressive and objective mathematical measures to judge creativity. Central to the measurement are two criteria – novelty and usefulness. The program is rated high on novelty if it manages to create something that is atypical of whatever is already fed to the program. There are various criteria for judging novelty pioneered by him. Besides novelty, the solution must be of high quality in order to be classified as creative. Also, emphasis on the quantity of high quality solutions is highly rated by him. The program must be capable of producing a substantial number of high quality solutions.

3.5 Relationship between Intelligence and Creativity

How is intelligence related to creativity? This question has no simple answer but many psychologists have drawn some empirical evidence to find the interrelationship of intelligence to creativity. The correlation between intelligence and creativity is positive up to approximately an IQ of 120. That is, an increase in intelligence leads to an increase in creativity up to the IQ of 120 (Simonton 1994). However, it is at this point that the correlation between intelligence and creativity breaks down. Once a person is at least above average intelligent, it is hard to predict creativity. Above a certain threshold of intelligence, exceptional creative achievement in science, art, music or writing ceases to depend on intelligence. Thus, psychologists have proposed: *Intelligence is a necessary but NOT sufficient condition for exceptional creativity.*

For exceptional creativity, intelligence is necessary but it is not the only requirement. There are other factors such as originality, perseverance, and risk-taking among many other traits that are more important once a certain threshold of intelligence is surpassed. However, this idea is not of much use to AI researchers because it was not deduced from a genuine understanding of intelligence and creativity. This idea was produced by observing humans operate in the real world. It was not produced by drawing upon the intrinsic working of intelligence and creativity. We argue in the next chapter that creativity has independent existence from intelligence in computers. This question of interrelationship between intelligence and creativity has not been adequately resolved by psychologists. For AI, we need to have a universal and independent understanding of intelligence free from human subjects. Just like Legg & Hutter (2006), we believe intelligence is a universal concept, that is, intelligence should not be understood to be unique to humans or anything else.

CHAPTER 4

INDEPENDENCE OF CREATIVITY

The important question to ask, as far as computers are concerned, is how does intelligence relates to creativity. Do we need to make computers genuinely intelligent before we make them creative? Does intelligence cause creativity or does creativity cause intelligence? Is creativity a process that can exist independently of intelligence in computers? Are they dependent on each other or are they independent processes? These questions have tremendous implications for general intelligence and creativity because the definite answers can prevent misdirected research efforts. If it is proven that one is the function of another, research in both Computational Creativity and Artificial Intelligence will receive tremendous forward progression.

4.1 Creative Programs

Many computer programs have been known to display substantial levels of creativity. AARON (Cohen 1994) is one of the programs that creates art. As of now, the art created by AARON can clearly be designated as human-competitive. Langley et al (1987) created several famous creative programs for scientific discovery. They claimed that their programs discovered Boyle's law, Ohm's law, Kepler's law and Coulomb's law among many others. Bringsjord and Ferrucci (2000) created the logic-based creative story-writing system called Brutus. Another historical program is Copycat (Hofstadter & Mitchell 1994), an analogy engine that discovers useful analogies in the domain of alphabets.

The above programs are creative but are they intelligent? In spite of their human-level creativity, they must not be considered intelligent because they work in very limited domains and obviously cannot stand any standardized test or definition of both human and machine intelligence. Hence, it can be argued that in computers, creativity can exist without intelligence. Below, we do a case study of Koza's Automated Invention system (Koza et. el. 1999, Bennett et. el. 1999) to demonstrate the independence of creativity from intelligence. However, our argument can equally well be justified from any of the other existing creativity programs. Koza's programs were selected for the sole reason that the author is familiar with them. His programs are just a representative of a class of all the creative programs that exist today.

4.2. Creative Unintelligent Systems – A Case Study of Creative Genetic Programming (GP)

Here we do a case study on Koza's (Koza et. el. 1999, Bennett et. el. 1999) systems known for creating automated inventions. He claims that his programs were able to generate many patentable inventions and hence convincingly display creativity. His GP based automatic inventions offered many useful insights into computational creativity and intelligence.

Genetic Programming is one of the members of the Evolutionary Algorithms family that evolves programs with mediums of variation and selection. We start with the initial population of programs and their selection is based on the fitness of individuals. The fitness function measures the individual's relative fitness in the population. An individual is a program (usually in the form of a tree structure) that breeds with other individuals through the medium of crossover to produce new offspring. Mutation causes small changes to particular individuals in the population.

Genetic Programming, as endorsed by Koza, has been remarkably successful in the process of discovery and invention. He strongly claims that his GP has been successful in producing human-competitive creativity in the domain of Analog Electronic circuit design. GP successfully infringed several existing patents and even generated new patentable inventions.

Most of the successes of GP have been in the domain of Analog Electronics. Analog circuit design is one of the domains where designers have few explicit rules and guiding principles. They need creativity, intuition and experience to contemplate a good design.

One of the areas of invention of Koza's GP has been electronic filters. Filters are used in electronic components to pass the desired part of the signal while eliminating the undesired signal. Only the desired frequencies are passed through to the next stage of processing. One of the classifications of the filters can be in the form of passbands – Lowpass, Highpass and Bandpass. Lowpass filters allow the entire signal below a certain cutoff frequency to pass through. High-pass filters allow the signal above a cutoff frequency to pass through. Bandpass filters allow a particular band of frequencies to pass through. Koza's programs have been used to synthesize most of the classes of these filters.

As an example, Cauer Elliptical Topology was generated by GP, without human intervention. A Cauer filter, invented by Wilhelm Cauer, was an advance over other filters of its time. Wilhelm Cauer was granted a successful patent for this work and it was classified as one of the breakthroughs of analog electronics. The filter has a sharp transition band; it can make a quick transition between stopband and passband. That is, it is true to the cutoff frequencies and very sharply attenuates an undesired signal. The GP successfully replicated the work of Wihelm Cauer without any human intervention.

The results can be considered creative because:

1. It was novel. The system generated the novel circuit that was not known to itself. It was non-obvious and had an element of surprise to it.
2. It was an appropriate filter. The filter sharply attenuates any undesired signal and passes the desired signal.

However, this program cannot be considered generally intelligent because:-

1. The system cannot stand any valid definition or test of intelligence, without exception. If we subject the system to the Turing test, IQ test or any other test of intelligence, it needs no genius to see that the system will fail.
2. It does not have the capacity to transfer or generalize the knowledge. It only works in a single finely-tuned domain.

Hence, on the above grounds, we conclude that Koza's GP is a creative unintelligent system. We consider Koza's experiments as an empirical proof of the existence of creativity without intelligence. This holds for many other creative programs developed over the years. On the contrary, attempts like CYC to demonstrate intelligence without creativity have failed. Even though the authors of the automated invention GP programs may themselves not agree with us; in this case study, we can safely conclude from their work:

In computers, creativity can exist without intelligence.

Creativity has been and will continue to exist in computers without intelligence. Creativity precedes intelligence; it's not the opposite as some psychologists will have us believe. The only generally intelligent entity, we know of at present, is humans, and they possess both intelligence and creativity simultaneously. No generally intelligent AI system has been successfully created to date but many artificial creative systems have been created. Creative systems can exist without general intelligence but there are no examples of vice-versa because it could be that creativity precedes intelligence. Hence, existence of general intelligence has never been shown to exist without creativity, neither in psychology, nor in AI, nor anywhere else. Now we are ready to make a bold claim that creativity leads to general intelligence in machines.

CHAPTER 5

A NEW DEFINITION AND HYPOTHESIS ON MACHINE INTELLIGENCE

Many definitions and theories of both human and machine intelligence emphasize Goal-Achievement. Consider some of these definitions:-

“Intelligence is the aggregate or global capacity of the individual to act purposefully, to think rationally and to deal effectively with his environment.” Wechsler (1944)

“*Intelligence is the computational part of the ability to achieve goals in the world. Varying kinds and degrees of intelligence occur in people, many animals and some machines.*” McCarthy (2004)

“*General Intelligence is the ability to achieve complex goals in complex environments*” B. Goertzel (Goertzel & Pennachin 2007)

After extracting the common features of several definitions of intelligence, Legg and Hutter (2007) proposed their informal definition of intelligence that describes the external working of an intelligent agent: *Intelligence measures an agent’s ability to achieve goals in a wide range of environments*.

Achieving goals in a wide variety of environments as a characteristic of intelligence can easily be questioned on many grounds. But this characteristic is theoretically and empirically useful in the context of present day AI research. The above definition tells us “what” an agent need to do (achieve goals) but fails to deliver “how.” The agent needs to succeed in multiple environments. Agreed. Now what? Where do we go from here? We need to know what an agent needs to do in order to succeed in multiple environments. We hypothesize that to achieve goals in a wide variety of environments; a system must be capable of being creative in a wide variety of environments. The *capacity for being creative in a wide variety of environments is a necessary and sufficient condition for General Intelligence in machines.* This idea was inspired by Marcus Hutter's AIXI model (Hutter 2004), the model that describes the behavior of an agent that behaves optimally in any computable environment. We intend to develop an empirical idea that allows the agent to be optimal and goal-achieving in a wide range of domains. Hence, we introduce our new definition of intelligence: *Intelligence is the capacity for being creative in a wide variety of environments.*

In defense of our hypothesis, we introduce three arguments: - the law of requisite variety argument, the complex systems argument, and the Turing test argument in the subsequent sections.

We reiterate that the term "creativity" here is not used in the sense of exceptional creativity. Creativity, as a term, is synonymous to normal creativity in humans, that is, the kind of creativity that exists in everyone. Normally, people associate creativity with highly creative individuals who make original contributions to a specific domain. That kind of creativity is rare but P-creativity is ubiquitous in everyday language, thoughts, actions and problem-solving.

It needs to be clarified that it is highly inappropriate to say that "creativity" is a necessary and sufficient condition for intelligence. Although creativity is necessary for intelligence (Dartnall 1994), the vagueness with the term creative in this context can create problems. As we saw in section 4.1, a program can be creative in one environment without displaying creativity in other environments. The word "creative" can easily be associated with the success of a system in one domain or a specific problem. But that does not say anything about its intelligence. The designer or programmer, by clever planning, can easily make a program that can achieve novel and appropriate behavior in a few environments. By specifying the plurality of environments, we ascertain that the system has complete control over itself, and any generally intelligent system is not supposed to depend on its designer.

Intelligence is a complex phenomenon and that complexity is not handled by the term "creativity" when it is applied in a single context**.** Learning, problem-solving, decision-making, reasoning, goal-achieving, and many more phenomenon are all known to be associated with intelligence. A system can be creative in one environment without possessing any of these. An Evolutionary Algorithm can easily attain creative behavior without possessing one or more of those attributes. Many such programs have been creative in art and music **(**Romero & Machado 2008). However, when the phrase "creative in a wide variety of environments" is used, it implies that the system will need all those attributes (learning, reasoning and so on) to achieve its goals.

Unlike existing creative programs, where the representations and design were hand-coded, to be classified as creative in a wide variety of environments, programs will need to display behavior that is classified as intelligent. What does being creative in a wide variety of environments imply? There are two points that can be made in this regard.

1.) The systems will need to create, change and design their own representations. No designer can hand-code all the representations because anticipating what kinds of environments a system will encounter is simply intractable. In order to display novel and appropriate behavior in an unknown environment, a system will first have to learn about that environment. It will need to form internal models of that environment (Holland 1995). Closely related is Conant and Ashby's (Conant and Ashby 1970) law of regulation "Every good regulator is a model of its environment." Hence the phenomenon of learning is implicit in our definition. Analogously, in the context of human-beings, it has been shown by researchers that before making any significant creative contribution, an individual has to learn a certain level of domain-specific knowledge.

2.) It can be easily argued that the definition does not include many of the crucial elements associated with intelligence – problem-solving, decision-making, dealing with uncertainty and so on. Moreover, there are many situations in which systems achieve their goals without being creative. The above criticisms are valid and sure to arise. Our response to these criticisms is that even though we agree that intelligence encompasses many elements, all these elements are implicit in our definition. There is a chance of misunderstanding here. So let's explain in detail what we mean. By saying that all the elements that are part of intelligence are implicit in our definition, we do not claim that all these elements (problem-solving, decision-making, information-processing, learning and so on) are synonymous with creativity (although the possibility is not ruled out). What we mean to say here is that if we can design a system that is capable of being creative in a wide variety of environments, it would be capable of everything else that is part of intelligence. If a machine is capable of producing both "novel" and "appropriate" output, then it is certainly capable of producing just "appropriate" output. Decision-making and problem-solving, even if they don't require novelty, they will require an appropriate output. The capacity for non-creative but appropriate output is obviously embedded in the capacity for creative output. Hence in our definition of intelligence, even the capacity for producing uncreative but correct output is embedded.

In the subsequent sections, we introduce three arguments in the defense of our proposition. The law of requisite variety argument reasons that in order to maintain the requisite variety to compensate for perturbations that hinder goal-achievement in each of the environments, a system must be creative in producing the requisite variety in the environment it finds itself in. The Turing test argument asserts that only a system capable of creativity across a wide range of domains can pass the Turing test. The complex systems argument reasons that in order to put up with the novelty of every complex environment, a system must be capable of creativity in each of those environments. All three arguments are not mutually exclusive; rather they all convey the same message but through a different line of attack.

5.1 The Law of Requisite Variety Argument

According to the law of requisite variety argument, in order to achieve its goals in a variety of environments, an agent must be capable of actively creating variety in its actions in a wide variety of environments. Each environment will generate its unique set of disturbances for an agent, therefore an agent must actively create corresponding responses for each environment. We argue that a generally intelligent agent must have the capacity to be true to the law of requisite variety and that is only possible if it is creative in multitudes of environments. In addition, we will argue that even in human-beings, evolution of general intelligence increased the variety in actions available in a wide array of environments. Imagine the following scenarios:

1.) Imagine a hypothetical robot that is not endowed with explicit rules to deal with every situation it might encounter. Now, we want this robot to succeed in game-playing environments. Suppose our robot is made a goal-keeper in soccer. In order to achieve the desired outcome (Goal defended) against the opposition team, it must create a new defense for each type of unique attack the opposing player might lodge.

2.) Suppose the new situation for this robot is to drive an automobile. In order to achieve the successful outcome (successful driving), it must create a new set of responses for every unique nuisance that the traffic creates.

3.) Imagine a generally intelligent system that is attacked by various unique computer viruses. If possible virus attacks are all unique and fundamentally different from each other, then a generally intelligent machine should have as many anti-virus responses as unique attacks. Because it is not possible to anticipate each of those unique kinds of virus attacks, our machine must actively create a response for such an attack.

But if humans also possess general intelligence, do they also possess requisite variety in the context of being creative in a wide variety of environments? We argue that the answer is unambiguously yes, and it will be apparent if we consider the evolution of general intelligence. Let's digress for a moment from a machine's case to a human case to see how general intelligence (popularly known as g) increases requisite variety in our species.

One of the mysteries of evolutionary psychology is how general intelligence evolved. Evolutionary psychology argues that the human mind evolved in a modular fashion where each module is dedicated to a purpose. Cosmides and Tooby (2002) proposed that dedicated and specific intelligences evolved to solve a particular evolutionary problem. Dedicated intelligences conferred reproductive and survival advantages on those who possessed them. Cosmides and Tooby (2002) assert:

"We humans solve many different adaptive problems well. To accomplish these feats, there must be at least as many independent evolved mental programs as there are adaptive domains in which the standards for successful behavior are qualitatively different. We think that one can identify hundreds or perhaps even thousands of these domains, ranging from thermoregulation, parenting, and food choice to mate choice, friendship maintenance, language acquisition, romantic love, pollutant avoidance, predator defense, sexual rivalry, status attainment, projectile accuracy, and kin welfare. Since environments cannot provide organisms with definitions of problem-solving success, independent problem solvers must be built in to the brain for each incommensurate value domain. For this and many other reasons, the brain must be composed of a large collection of evolved circuits, with different circuits specialized for solving different problems. In this view, the brain is necessarily a diverse collection of dedicated computers networked together."

Each of these modular intelligences is a unique response to a specific perturbation in the environment. These modular intelligences, by solving a specific evolutionary problem, increased the variety of actions available to our species in the face of each unique challenge. If dedicated or specific intelligences served the purpose of evolution, why would general intelligence evolve? This enigma of evolution of general intelligence is one of the most important problems in evolutionary psychology.

One of the answers is provided by Cosmides and Tooby (2002). They draw distinction between dedicated intelligence and improvisational intelligence (their term for general intelligence). Dedicated intelligence, as said earlier, is a specialized form of adaptation in response to one kind of evolutionary challenge. However, no matter how many dedicated intelligences and modules the human brain evolves, there will always be problems that are novel and out of range of those dedicated intelligences. Improvisational intelligence, in the face of novelty, is nothing but the coming together of different dedicated intelligences to solve a novel problem. However, when these dedicated components act together, their behavior is typical of complex systems, and they respond to novelty by creating novelty of their own. While we are writing this research paper, it is our improvisational intelligence that is called upon. In the modern world, most high-status professions require high improvisational intelligence or more popularly, high g (Gottfredson 2008).

Another competing but controversial theory is of Kanazawa (2004). He proposed that the evolution of general intelligence is a domain-specific adaptation to solve problems that are evolutionarily novel; problems which are new in the sense that there is not enough time for evolution to create a separate dedicated module. Further, he argues that the problems that are not evolutionarily novel confer no advantage to highly intelligent individuals. According to him, domains like parenting, dating, mating and cheater-detection are non-novel evolutionary problems; hence individuals with high g don't do any better than those with low g. But domains like mathematics, philosophy, engineering and administration are examples of evolutionarily novel problems and as expected high g individuals excel at these whereas low g individuals are challenged with these problems. Unlike Cosmides and Tooby's theory, this view asserts that general intelligence has separate existence from dedicated intelligences.

Whichever view may be right, both of these views imply that general intelligence, once evolved, increased the variety of actions available to our species. Surprisingly, the creation of requisite variety is apparent from both perspectives - humans as a collective species and human being as an individual. Human-beings, as a collective, have historically been successful in creating variety in a wide variety of environments. They survived predators, found cures for fatal diseases, invented the wheel, found ways to escape natural calamities, and most importantly, competed successfully against other groups of their own species. If we trace back through the past, the problems that have been solved by humans throughout their evolutionary history have required creativity in a wide variety of contexts. Whether it be the creation of tools in the hunter-gatherer era, creation of agriculture, creation of medications, creation of literature and language, creation of religion, or creation of science, it all required tremendous creativity in a wide variety of contexts. The world is now much more complex than it has been in the entire evolutionary history, such complexity puts far more demands on creativity. Again human beings, as a collective, are true to the law of requisite variety by creating responses for each unique challenge in each unique environment.

Human beings rely on other humans for their survival. Most of the discoveries and achievements, crucial to survival, can only be credited to the few. The evolutionary success of the other members has always depended on the crucial work of those few. It is precisely for this reason that it is hard to see the *capacity for being creative in a wide variety of environments present in human-beings* but it has always existed when we see "human-beings" as a collective.

In addition, human beings are true to the law of requisite variety even at the individual level. When we confront novel problems in routine activities, we create new responses that are appropriate to the situation. Thinking of a gift to give to a date, coming up with a new cooking recipe, a student inventing an excuse to turn-in a late assignment, writing a funny message on facebook, and something as simple as the routine use of language (Chomsky 1966) require us to create variety in the face of disturbances in daily life.

However, machines do not need to play an evolutionary game like us. Whether we will have a general intelligence in a single machine or a "collective species of machines" is not an easy question, but here we use the word machine to incorporate both cases. If a single generally intelligent machine is created, it must single-handedly create variety in its actions for each corresponding environment. If a "collective species of machines" is created, the collection as a whole must have the requisite variety in its actions. The complexity of the world now is higher than the complexity throughout our evolutionary history. A machine must be able to deal with ever-increasing complexity and this is only possible by actively creating the variety in its responses to any challenges posed by the environment.

Let's consider set-theoretic terminology similar to Ashby (Ashby 1958). Assume D is a set of disturbances and R is the set of responses. Now assume in a new environment that a new disturbance d_i pops up and that this d_i is significantly different from the elements of disturbance set D which the agent has already encountered. None of the responses from R can counteract the disturbance because d_i is significantly different from already existing disturbances. To mitigate this new disturbance, a system must be able to create a new response r_i, which is significantly different from the elements of existing set R. This new response can be said to be creative because it is novel ($r_i \notin R$) and appropriate (able to counteract d_i).

Now this capacity of a system to create a new response in one environment does not automatically get transferred to other unrelated environments. Many creative programs (as we saw in section 4.1) are able to display creativity in only one domain. In one environment, the elements in the set of disturbances D may not be very different from each other; so it is possible that the elements of the corresponding response set that the system creates may not be very different from each other. As a result, the system may display creativity in a single environment but there is no guarantee that it can display it in many environments. Just like most of the existing creativity programs, its creativity may be confined to a very limited domain; contrary to our goals. Ashby summed it up in this quotation "Change the environment to its opposite and every piece of wisdom becomes the worst of folly."

Hence this capacity for being creative must be present in the system across a wide range of environments. An ideal agent is one that is capable of creating new responses to the demands of a unique environment. To maintain the requisite variety in its actions for each environment it encounters, an agent must be able to display creativity in each environment. That is, the requisite variety is actively created by the system; it is not passively present in an intelligent agent.

5.2. Turing test Argument

To pass the Turing test, a machine obviously needs a linguistic ability on the order of human beings. Human beings are creative in their use of language - Chomsky calls this phenomenon the Creative Aspect of Language Use or, popularly, CALU (Chomsky 1966). The CALU argues that humans are very effective at generating novel and appropriate sentences. While writing an article or conversing with another human, human-beings bring together words in a unique way that is appropriate to the situation. Chomsky, in his CALU, asserted that it is quite remarkable that using only a few rules of grammar, humans are able to generate such a rich variety of sentences. An exact similar capacity must be possessed by machines, if it is to fool a human judge on the Turing test. Chomsky (1997), while referring to "Descartes Problem" stated that it is only by possessing CALU, that any creature can be credited to possess a mind. Even though our focus is not on mind in this thesis but according to Chomskian arguments, any machine, if it is incapable of CALU (or simply creativity), cannot pass the Turing test. Even Turing believed that only a machine capable of creativity can pass his test, as is evident from his answer to the Lovelace objection (Turing 1950). Creativity is required for passing the Turing test; at the minimum linguistic creativity is necessary. Consider this sentence that I spoke:

I went to the University's recreation center at 9:40 AM on July 4 but it was closed, so I decided to go swimming at my friend's pool.

Though this sentence looks very ordinary, nevertheless, it is creative. This sentence is new; an exactly similar sentence was never been generated before; at least in the context of this speaker. To pass the Turing test, a machine will need to possess this linguistic creativity similar to humans. The machine will need to generate sentences that it has never seen before, never spoken before, never heard before in order to pass the test; hence demonstrating creativity in the P-creative sense.

Moreover, the questions can come from a variety of domains, and a good judge will not test the machine on questions that are very similar to each other. He will test the machine on a wide variety of unrelated questions. Hence, *only the machine that is capable of generating creative responses in a variety of situations will pass the Turing test.*

5.3. Complex Systems Argument

Godfrey-Smith (2002), in his Environmental Complexity Thesis, suggested that the purpose of cognition is to deal with complexity. Why do we care about learning, judgment, decision-making, intelligence, creativity, problem-solving, reasoning and everything else that is cognition? It is only because the environments are complex and we need all those fancy capabilities to handle the complexity.

Complexity researchers consider the term “Perpetual Novelty” as the characteristic of every Complex Adaptive System (CAS), which are systems that contain many components in interaction with each other resulting in global emergent behavior. Ecosystems, ant colonies, immune systems, national economies, organizational structures, genetic algorithms and human brains are all known to entertain perpetual novelty. But this novelty does not magically appear in the environment. It is not created out of nowhere. Rather, it is the result of the already existing building blocks (Holland 1995). Even though the building blocks are already present in the environment, novelty emanates from the recombination of these building blocks (Holland 1995). Without the danger of over-simplifying, we can safely state that it is the permutations, combinations, blending and amalgamation of building blocks that contribute to novel behavior in the environment. But when we talk of building blocks unifying, we must not forget that the unified building block is “more than the sum of its parts.” That is, the properties of the resulting building blocks are non-linear; it is not possible to sum the behaviors of components to predict the behavior of the whole. Also, the behavior of the whole will be emergent or in other words, there is a disconnect between global and local behavior. Hence, our agent will find itself in situations composed of already existing building blocks but which are new in the sense in which these building blocks come together. Holland’s these ideas imply that even though the challenges posed will be somewhat novel, there will always be some order in the environment. The entire crux of the above argument is that an agent will necessarily encounter complex environments and new situations, yet with some pattern to them.

What if the system has not been trained for such new environments and situations? What if the system does not have a method for handling such situations. What if it doesn't have explicit resources to deal with novelty? What does an intelligent agent do now? It has encountered a new situation. It must collapse. But no, there is a way out. According to Holland (1995):

"This provision for simultaneously active rules helps us understand an agent's ability to handle a perpetually novel world. It contrasts sharply wherein the agent has only a single rule for each situation. With the single-rule approach, the agent must have rules for every situation it may plausibly encounter. This poses a problem analogous to the one we discussed earlier for the immune system. An agent is unlikely to have a single rule adequate for each situation it encounters for the same reason that the immune system lacks a set of antibodies prepared ab intio for all possible invading antigens-there are just too many possibilities. With simultaneously active rules, the agent can combine tested rules to describe a novel situation. The rules become building blocks.

By way of example, consider someone in the unfortunate circumstance of having a "flat tire while driving a red Saab on the expressway." Most of us have not driven a Saab, let alone had a flat tire while driving one, but we would not be at a loss for an appropriate response. The reason would seem to be that we decompose the situation into familiar parts. Most of us have had some experience with flat tires, or at least know procedures for dealing with them."

What he is saying is that a CAS agent can combine internal rules to deal with a novel situation. Assume an agent knows how to deal with situations A and B separately. Now, both of these situations occur simultaneously. It would be inappropriate for an agent to use rules that it learned to deal with either of A or B. The agent, at this point, in order to achieve a successful outcome, must create a new rule. This point will be apparent if we consider the following example. Suppose a designer trains a robot to move backward when it detects any moving entity from the front and move forward when it detects a moving entity from the back. What if two moving entities approach from both front and back? What must the robot do in order to protect itself? It can't use the rules it knows because both rules coded by the designer lead to failure. Hence it must display behavior that is new in this situation.

The system must cash in on its limited resources to extrapolate a new strategy and thereafter succeed in it. To succeed in such unknown conditions, the system will have to discover novel strategies that were not already known to it. Clearly, all that the system needs is to display creative behavior, that is, the behavior that is novel and appropriate while operating under computational constraints. Goal-Achievement in the complex environment depends on the agent's ability to successfully comprehend its environment and optimize its resources to comprehend the environment. If an agent encounters a new situation it has not been trained for, it must display creativity for the obvious reasons that it already does not have enough mechanisms for dealing with such novelty.

According to this Complex Systems Argument, the only way for an agent to achieve goals in a variety of complex environments is by being creative across all environments. Each complex environment poses novel challenges for an agent. To deal with these novel challenges, an agent will need to create new solutions and rules for each environment. Holland (1995) suggests that since it is impossible for an agent to have rules for each possible novel challenge it might encounter, an agent must create new rules by blending or combining internal rules. Thus, an agent, by being creative, can overcome the novelty imposed by a wide range of complex environments.

5.5. DERIVATION OF OUR DEFINITION FROM AI RESEARCHERS' DEFINITIONS

In this section we show how our definition is a natural consequence of other important definitions introduced by AI researchers. The following definitions are taken from Legg and Hutter's (Legg & Hutter 2007) repository of AI researcher definitions.

Consider the following two definitions:-

"…the ability of a system to act appropriately in an uncertain environment, where appropriate action is that which increases the probability of success, and success is the achievement of behavioral subgoals that support the system's ultimate goal." J. S. Albus

"Intelligent systems are expected to work, and work well, in many different environments. Their property of intelligence allows them to maximize the probability of success even if full knowledge of the situation is not available. Functioning of intelligent systems cannot be considered separately from the environment and the concrete situation including the goal." R. R. Gudwin

The only way to "maximize the probability of success even if full knowledge of the situation is not available" is to come up with many creative solutions for a given problem in a given situation. If an agent does not come up with creative solutions, it is hard to see that it can maximize its probability of success. An uncreative agent will be significantly limited in its ability to search as exhaustively as a creative agent. It may perform well in one environment, but it will need to rearrange its elements and representations in order to comprehend an unknown environment making creativity a necessary component.

"Any system …that generates adaptive behviour to meet goals in a range of environments can be said to be intelligent." D. Fogel

Generating adaptive behavior to achieve goals requires creating novelty in the direction of the goal. Hence, the only way a system can meet goals in a variety of environments is by being creative in those environments. Each environment poses new challenges and only the system capable of handling novelty will survive the challenges of that particular environment.

"Achieving complex goals in complex environments" B. Goertzel

"*We define two perspectives on artificial system intelligence: (1) native intelligence, expressed in the specified complexity inherent in the information content of the system, and (2) performance intelligence, expressed in the successful (i.e., goal-achieving) performance of the system in a complicated environment."* J. A. Horst

There are three processes necessary in order to accomplish complex goals in complex environments: to comprehend the complex environment, to shape the environment and to adapt to the environment. An agent will have to comprehend its environment. This process of comprehending the complex environment will itself call for creativity because the agent will have to develop new representations and understandings that previously did not exist for it. Secondly, the process of generating output in the environment itself leads to creativity because the environment will be shaped in a way that is both novel and appropriate with respect to an agent. A similar idea has been expressed in Robert Sternberg's theory of Successful Intelligence (Sternberg 1999) in the context of human-beings. Regarding adaptation, there is no shortage of literature linking adaptation to novelty and the creative process (Holland 1995, Holland 1975).

"Intelligence is the ability to use optimally limited resources - including time - to achieve goals." R. Kurzweil

"Intelligence is the ability for an information processing agent to adapt to its environment with insufficient knowledge and resources." P. Wang

The idea that an agent is limited in resources itself implies that an agent will need to best optimize those resources. In an uncertain environment, the system's success depends on its ability to make the best possible use of its resources. Either new uses for the resources must be found or the resources must be combined in a new way so as to achieve goals.

"...doing well at a broad range of tasks is an empirical definition of `intelligence'" H. Masum

A generally intelligent agent, by its very definition, must not be specifically trained for a broad range of tasks. If it is specifically trained for some task, it will not be capable of performing on a newer task. It will be necessary to "transfer the knowledge of one domain to another", a process very commonly associated with creativity.

"...in any real situation behavior appropriate to the ends of the system and adaptive to the demands of the environment can occur, within some limits of speed and complexity." A. Newell and H. A. Simon

Complex environments are chaotic, dynamic and non-linear. The "demands of the environment" are constantly changing and only the system that is creative can achieve "behavior appropriate to the ends of the system" in a dynamic environment. If there is an unknown transition in the state of the environment such that the system has not been trained to handle it, the only way for a system to succeed is by coming up with novel strategies. Limits of speed and complexity further necessitate that a system be capable of coming up with creative solutions.

CHAPTER 6

A NEW TEST FOR MACHINE INTELLIGENCE

Grounded in our definition and hypothesis, as introduced in the previous chapters, we introduce a new test for machine intelligence. This new test emphasizes c*reativity in a wide range of domains* as the performance criteria for machines. To reiterate what we introduced in chapter two, creativity, as discussed in this paper, is routine creativity rather than exceptional creativity. Hofstadter (1985) contends that both genius-level creativity and normal creativity emanate from the same underlying process; what is different is the matter of degrees not processes. Hence, creativity, as discussed in this chapter is in the normal P-creative sense.

6.1. Lovelace Objection

One of the major objections to the Turing test has been from Lady Lovelace (Turing 1950):

"The Analytical Engine has no pretensions to originate anything. It can do whatever we know how to order it to perform."

However, the Lovelace objection that computers will do whatever we program them to do and cannot go beyond what they know, has already been empirically proven wrong by many creative programs, as discussed in Section 4.2. Evolutionary algorithms are inherently unpredictable. Even though we believe that the Lovelace objection is invalid, one of the constructive parts of the objection is its emphasis on creativity. Any real test of machine intelligence must have creativity at its core.

In order to address the issue of creativity, Bringsjord et. el. (2001) introduced the Lovelace Test. According to them:

The Agent A, designed by designer D, passes the Lovelace Test if

1. A generates output such that this output is not the result of some coincidence but of processes that are repeatable and
2. The designer D could not have predicted the output from A's components (both hardware and software).

The Lovelace Test incorporated creativity but there are some major problems with it. Any machine that can create an output O even when it is only in a single domain, such that the designer could not have predicted the outcome, the machine will be attributed successful according to this test. If our algorithm regularly creates interesting music that surprises us, should we deem it intelligent? If a complex evolutionary algorithm is successful in solving some problem, such that it surprises its designers, then it must be designated successful according to the Lovelace Test. However, one of the constructive parts of their test was its emphasis on creativity.

6.2. Requirements for a Good Test

Grounded in our definition of intelligence as a function of creativity, we propose a new test. If we are to construct a test based on our definition, what must the test include? In the next few sections, we describe the factors involved in such a test.

6.2.1. Capacity for being creative in n distinct domains where n is a sufficiently large number

An agent can only be considered intelligent if it has the capacity for displaying creative behavior in a sufficiently large number of environments. In addition, these environments must not be too close or too similar to each other; otherwise there is a chance that an agent's success is just confined to similar environments, contrary to our definition. To make sure that the domains are dissimilar, either a human judge must be used or mathematical measures must be invented. Our test does not place any constraints on the subjectivity or objectivity of how distant the domains must be. It is possible to make decisions in the test both by a human judge as well as rigorous measures, which may not be hard to invent as we saw in section 3.4.

How large should be the n? An agent that is creative in three domains is *likely* to be more intelligent than the one that is creative in a couple of domains. But an agent that is creative in ten domains has a higher probability of being more intelligent than the one that is creative only in a couple of domains. If one agent outperforms another in creativity in many domains, then it is better than the other. Once again, we do not place any constraints on the issue of subjectivity or objectivity of the value of n. The decision could be made by a human judge. On the other hand, some objective value of n should not be hard to invent by designers of generally intelligent machines.

6.2.2 Able to replicate creativity in the domain under consideration

The outputs generated in a particular environment must not be the result of a hardware fluke (Bringsjord et. el. 2001) or random chance. The system must be capable of replicating creativity for each domain under consideration. Here it is important to mention that it is not necessary for the system to replicate the exact output in a particular domain; rather it must be capable of replicating creativity in that domain. That is, even though the output may be different everytime, it must be creative at least most of the times.

6.2.3 Intrinsic creativity

Holland (1995) introduced an idea of internal models:

"*the models of interest here are interior to the agent, the agent must select patterns in the torrent of input it receives and then must convert those patterns into changes in its internal structure. Finally, the changes in structure, the model, must enable the agent to anticipate the consequences that follow when that pattern (or one like it) is again encountered….*"

The internal model aids an agent to anticipate its environment so that appropriate action can be initiated. As the environments change or an agent's information content about the environment changes, an agent needs to generate new internal models in that environment; hence it must be creative intrinsically. The internal representations will also need to change in order to adjust to the new situations. To accurately assess the environment, novel plausible hypotheses must be generated and confirmed, so that positive reinforcement from the environment can be obtained. Even perception of the environment may not be straightforward for a complex agent in a complex environment. There may be a plethora of possibilities even in the interpretation of the environment. Hence, the system must be capable of intrinsic creativity to comprehend its environment.

Intrinsic creativity can be said to be associated with the processes that are not apparent in the agent's environment. It has to do with the internal working of an agent. An agent can be said to be intrinsically creative when it displays creativity in the processes (or whatever) that are internal to it. These processes may not directly affect the agent's behavior or output in its environment but help the agent anticipate and succeed in its environment.

6.3. Creativity–Based Intelligence Test (CBIT)

Finally, we have reached a point where we can describe our new machine intelligence test. To summarize our test, an agent must be capable of creativity in a variety of domains such that the creativity in each domain may be replicated. In addition, it must be capable of intrinsic creativity. Hence, the Creativity–Based Intelligence Test:

Definition : An Agent A is considered intelligent, if it is

(i) able to produce Outputs in n distinct domains where

a) n is a sufficiently large number and

b) the Outputs are considered creative in the domain under consideration

(ii) able to replicate creativity (not necessarily a particular output) in the domain under consideration

(iii) able to display intrinsic creativity in each domain.

The test can be conducted at three levels; at one level we compare the “creativity in a wide variety of contexts” of the machine to the creativity of human-beings as a collective, at another level we compare it to the capacity of a single individual, and at the third level a machine can be compared to another machine. Which level is the most appropriate can be decided by the designer for her own specific project.

At the level where a machine is compared to a collective species of human-beings, a machine that is creative in areas such as verbal behavior, music, mobility, art, sports, science and so on can be commended to have passed the test. However, at the individual human intelligence level, it must be able to display normal levels of creativity in many domains. If it is able to generate new and appropriate sentences in the verbal domain, able to generate new and appropriate movements in the “mobile robotics” domain, able to solve routine mathematical problems creatively, able to create a new strategy in chess, while changing its internal models in the process, then we can say it passes our test. At the third level when a machine is compared to another machine, a machine that can display creativity in more domains is more intelligent than the other.

6.4. Who Judges Creativity?

The important question is “Who judges the creativity of machines?” Are human judges appropriate? Our answer is yes until we have good objective measures of a program's creativity. There are many tests that require a human judge (like the Turing test and its derivatives). The question of subjectivity may arise but we see that in the real world humans are not bad judges of creativity. In almost every field, exceptionally creative individuals are judged by human referees of that field. But in our test, judges need to test routine creativity, which should be comparatively simpler. Some researchers may point out that this element of subjectivity may make our test less rigorous and informal. We argue that even if that is the case, it does not make our test weaker, but stronger. Designers of intelligent agents are humans, at least at present. Hence by applying their own criteria of creativity, they can make informal judgments about their systems at various stages of completion.

Nevertheless, this requirement of human judges is not a necessity. We already saw in section 3.4 that work on formal measures of computational creativity have been convincingly pioneered by Ritchie (2001, 2007). Objective measures of a program's creativity are soon going to be invented. Thus, our test has both subjective and objective character.

6.5. Predictions

1.) The systems that measure high on any general intelligence measure will also measure high on creativity (on any valid scale). If it is shown otherwise, then our hypothesis is invalid. This hypothesis is in the context of machines only. It is necessary to specify here that the term "intelligent" can have many varied meanings. Many of the classic programs like Mycin can be regarded as intelligent but not "generally intelligent." Hence, it is "General Intelligence" specifically that depends on creativity.

2.) Any machine that is incapable of creativity in a variety of situations will never be able to display human-like intelligence or any other kind of general intelligence. Dartnall (1994) considers it axiomatic that if machines cannot create, then they cannot be intelligent. Our hypothesis implies that the uncreative systems like Cyc must fail. If Cyc succeeds in creating common-sense reasoning or general intelligence, then our hypothesis is wrong. There is near consensus among researchers now that Cyc has already failed and its results are very disappointing.

6.6. Objections

Due to the nature of the topic, we expect objections. Here we answer some of the potential objections that can be raised against our work:-

1. This definition and test cannot stand Searle's Chinese room argument.

What we proposed here is a working and simple definition of intelligence that has empirical significance. As far as the "understanding" and "mind" go, addressing these issues was not the concern of our paper. We assume that it is possible for systems to possess intelligence without possessing any of these. As discussed in chapter 5, it is implicit in our definition that if a system is creative in a wide variety of domains, then it will be capable of adaptation, goal-achievement and problem-solving in those domains. But we don't think it is appropriate to address the issue of mind and understanding for empirical work at present.

2. The program's creativity is difficult to measure. The criteria for judging creativity can be very subjective.

Some authors (Ritchie 2001, Ritchie 2007) have already shown some impressive work in the assessment of creativity. The assessment of creativity, well documented in their work, is not hard to quantify in mathematical terms and good objective functions can be developed. However, we do not have a requirement that creativity be self-assessed by the program. Even human judges would serve the purpose for our test. Human judges are used in many of the machine intelligence tests (like the Turing test).

3. Why is the creativity definition better than definitions stressing Goal-Achievement of an agent?

The Goal- Achievement definitions introduced by many authors tell us “what” but do not tell us “how.” Our definition is more generic and tells us how an agent needs to achieve goals (by being creative) in a wide array of domains. In complex environments, there is no way an agent can achieve goals without being creative.

4. Is not creativity a different facet of intelligence? Why should we believe that creativity causes intelligence?

If a machine is capable of producing both novel and appropriate output, then it is certainly capable of producing just “appropriate” output. Decision-making and problem-solving, even if they do not require novelty, require an appropriate output. The capacity for non-creative but appropriate output is embedded in the capacity for creative output. If a system is capable of creativity, then it is also capable of problem-solving that does not require creativity.

5. Why generally intelligent systems cannot be created without creativity?

It turns out to be the case that any attempt to create a highly intelligent system without creativity has not been fruitful. In the last fifty years, researchers made this assumption, but no uncreative intelligent system has been created to date. There are systems like Mycin, Dendral and several others that do extremely well in narrow domains but they are not generally intelligent. Their intelligence and performance is confined to very specific problems. The failure of Cyc is the grandest example of an attempt to create an uncreative intelligent system. Cyc tried to create a system based in predicate calculus that could not generate novel behavior. Cyc’s failure remains the empirical proof of the non-existence of intelligence without creativity.

Koza and colleagues (Bennett et. el. 1999) suggested that illogical thinking plays a significant role in creativity. However, he was not the first person to suggest this; many creativity researchers have known the value of illogical processes for a long time. Pure deductive logic or systems that are too tight in reasoning overlook some of the best solutions. Some degree of randomness is necessary for escaping the constraints of logic. Something analogous to mutation in Genetic Algorithms is always a part of intelligence and creativity. We end this chapter in John Koza and his colleagues' words "logic considered harmful."

CHAPTER 7

DESIGNING INTELLIGENCE AS A COMPLEX ADAPTIVE SYSTEM

In this thesis, we proposed that in order to display general intelligence, a machine must be capable of creativity in a wide variety of environments. This chapter briefly asserts that in order to attain creativity in a wide array of domains, the way knowledge is managed in a system must be a complex adaptive system. Hence, the chapter argues that any knowledge based system that aims to create general intelligence must be a complex adaptive system. There are three mechanisms that must be part of such a knowledge based system, viz., interdependency and fluidity, mechanisms for attribution of emergent properties and self-organization.

7.1. Introduction

Although the time has come that our perception of the whole field of AI should be of complex systems and more rule-based, well-defined and logic style approaches need rethinking, we confine ourselves to the discussion of Knowledge-Based Systems in this paper. Loosemore (2007) argues

"One the most basic assumptions made by Artificial Intelligence researchers is that the overall behavior of an AI system is related in a lawful, comprehensible way to the low level mechanisms that drive the system.……..this apparently innocent assumption is broken, because all intelligent systems, regardless of how they are designed, must be complex systems"

Evidence in anthropology and psychology suggests that mental fluidity is central to human cognition. It is this fluidity that lends us our higher level abilities that only our species Homo sapiens sapiens possess. If our goal is to build general intelligence, the only valid path is complex adaptive systems – since by their very nature these systems can efficiently produce an analog of mental fluidity. Human cognition is a complex adaptive system (Morowitz and Singer 1995), hence KBS have to be complex adaptive, it cannot be otherwise. For the purpose of this paper when we talk of KBS, it encompasses all areas of Artificial Intelligence that handle knowledge including case-based reasoning, expert systems, and so on. Complex systems are the systems that involve interactions among various components and as a result of these interactions global emergent behavior emerges in a system. Bar-Yam (1997) suggests "interdependence" among various components is a more generic term since it is the interdependent parts that create emergent properties rather than interconnected parts. Complex systems are known for various characteristics; among them widely discussed are emergence, self-organization and non-linearity. Here it is important to mention that KBS must be complex adaptive rather than just complex because they must be sensitive to any new knowledge that enters the system.

There are three mechanisms that must be part of any knowledge based system, viz., interdependency and fluidity, a mechanism for attribution of emergent properties and self-organization. These mechanisms will not only make knowledge more complex, but also make it adaptive to new information that comes in.

7.2. Interdependency in Knowledge Bases

Referring to the brains of early humans, Mithen (1996) asserts:-

"we can safely state that in spite of linguistic differences, all Early Humans shared the same basic type of mind:a swiss-army-knife mentality. They had multiple intelligences, each dedicated to a specific domain of behavior, with very little interaction between them……..Early Humans seem to have been so much like us in some respects, because they had these specialized cognitive domains; but they seem so different because they lacked a vital ingredient of the modern mind: cognitive fluidity"

Grounding his arguments with evidence from archaeology and anthropology he further suggests that "cognitive fluidity" is responsible for our intelligence and the rise of human civilization (including art, science and religion). Even our closest cousins Neanderthals did not have this ability. But it is very unfortunate that AI researchers and builders of KBS have totally disregarded this ability and created systems that stored information like Early Humans. Cyc

(Lenat & Guha 1990) did better by at least creating Neanderthal-like common sense, that is, there are not very flexible across-the-domain connections. The time has come that our perspective changes and we start treating knowledge as complex for AI purposes. Unfortunately what should have been done first will be done last. All the KBS in the past, without exception, were fundamentally flawed because they represented knowledge "as is" without giving any consideration to how a thing can be represented in terms of other entities.

Besides anthropology, there are theories by psychologists and computer scientists to account for our mental fluidity. Arthur Koestler (1964) proposed the Bisociation, the mechanism in the human mind, through which different planes of thought come together to produce a novel thought. Karmiloff-Smith (1992) came up with a Representational Redescripton model explaining how children's representations become more flexible with age. Fauconnier and Turner (2002) proposed the theory of conceptual blending, in which multiple concepts can blend to form a new concept with emergent properties. Boden (1991) also suggested that transformation of conceptual spaces is central to our thought processes. Hofstadter & Mitchell (1994) came up with a program called Copycat to model analogical thought but they asserted it can be extended to all fluid mental concepts. The major achievement (Mitchell 2001) of the Copycat program was to show that human cognition is a "complex adaptive system." The above works unanimously offer a unique insight into the workings of the human mind - In the human mind, elements from different areas are interdependent and often come together to form a coherent whole whose properties may be emergent.

Creating highly fluid systems that are complex and adaptive is what is required. KBS should be capable of "Seeing one thing in the frame of another" and free merger of concepts should be the rule. High-Interdependency among representations must be the rule in any knowledge base. Each concept must be represented from the perspective of other concepts in the knowledge base. And a concept should have representation from the perspective of multiple other concepts. This is done to ensure high interconnectivity, which obviously will not only make KBS highly fluid but also adaptive, just like human cognition. French (1995, 1998) has argued that representation for any AI system, if it is not malleable, then it is necessarily flawed. According to him, the representations that are fixed can never produce an analog of human cognition and fluidity. Each representation should not be "as is", that is, there must not be a fixed representation of any concept.

7.3. Emergence

The reason emergence becomes a central issue in the design of knowledge bases is that there are global properties possessed by entities which are not reducible to the components and subparts of the entities. For example – psychological characteristics of the human brain cannot be explained in terms of just one neuron or many neurons taken independently. Rather interconnections and interdependence of the neurons display emergent behavior.

Let us consider a hypothetical knowledge base. Assume the representation to be either predicate calculus, logic, semantic network, frame based or any symbol manipulation representation. Also, we assume our knowledge base to be all-inclusive, that is, containing all the information in the world. Suppose we examine the concept of "Automobile" in our knowledge base. If we start with any possible initial state of the concept "Automobile" in any representation, and from there try to derive this property "Means of transportation", we see that no matter what change we make to that state it is simply impossible to derive the property "Means of transportation" bottom-up since derivation is the process of sequential state transition in which a part or subpart of the information is changed to reach an end state. However, the emergent property can never be an end state because:

1.) The concept "Automobile" is represented in the form of its parts (engine, tires, windows, mirror, doors and their relationships). Information "Means of transportation" is not contained in any one part, or combination of a few parts, rather it is the interactions of all the parts that emerge this property, hence, trying to extract information "Means of transportation" from the parts is impossible. This property "Means of transportation" is the property of a specific configuration created by all of the components and their interactions making up the concept "Automobile". Derivations, as in any representation lead to an end state in a deterministic fashion, however starting from any initial state in this problem we can never be sure that the next state will lead us to correct end state (our property). Hence the property is not derivable from the inner components and is globally emergent.

2.) To derive the property from all the other information in the knowledge base besides the "Automobile" concept, it is that the property if it is the property of another entity in the knowledge base, then it is going to be an emergent property of that entity, that is, not derivable from its own components. For example - the property "Means of transportation" can be attributed to the concept of "Horse" and it is non-derivable for this concept. Since we have already seen that the property cannot be obtained from its own components, the only way is for the property to transfer from one point in the knowledge base to another point by transferring from one whole to another whole. The property of the concept "Horse" can be transferred to the concept "Automobile" as a whole.

We need to have such mechanisms in our systems that help us in attribution of emergent properties because trying to derive any higher order property from lower level components is almost impossible. Analogical reasoning is one such mechanism but there are many more similar strategies that must be incorporated in the systems. The idea is to transfer concepts from one whole to another based on some criteria, and not try to derive things bottom up.

7.4. Self-Organization

Karmiloff-Smith (1992) while proposing her RR model wrote

"My claim is that a specifically human way to gain knowledge is for the mind to exploit internally the information that it has already stored……..by redescribing its representations or, more precisely, by iteratively re-representing in different representational formats what its internal representations represent."

The above lines clearly hint to the process of self-organization in the human mind. A process analogous to the re-representation mechanism is a necessity in a complex adaptive KBS. Langen et. el. (2004) suggested self-organization is a requirement for any system if it is to be creative. The idea of incorporating self-organization is simple – any new information is connected to other information and this interdependency can lead to potential effects on all the interdependent information. The goal of self-organization is to adapt the whole KBS to any incoming information so that all the information is represented in the most accurate state in the KBS. The accurate state in the KBS is variable since induction of any new information can change the accuracy. For example – in the nineteenth century, the most accurate explanation for the universe was in terms of Newtonian physics, however, in the twentieth century after new information came in (Theory of relativity), the most accurate explanation for the universe was in terms of relativity.

The more inter-dependent various fragments of knowledge inside the knowledge base are, the more they are prone to the effects of the induction of new knowledge. This new knowledge, if it modifies any knowledge in the KBS, can subsequently lead to a chain of modifications because this modified knowledge is highly interdependent with various other fragments of knowledge. The system must be capable of self-organizing at this stage, that is, just by using all the information that is internal to the KBS, the system should be able to reach the most accurate representation or perception or frame for each fragment of knowledge that is effected by newly induced knowledge in some way.

7.5. Conclusion

No doubt, the road to making knowledge complex and adaptive for AI is filled with some serious bottlenecks, nevertheless, it is reachable. Designing complex adaptive KBS is the most ideal approach that will make them adaptive by incurring several advantages over conventional systems. The most significant advantage being an increase in information resulting from high interdependency in the knowledge base. Since each entity or situation can be perceived in multiple frames, our systems will have an option to choose the best frame that is the most accurate representation relative to other representations, thus increasing the system's information, reliability and adaptiveness.

REFERENCES

Ashby, W. R. (1958). Requisite Variety and its implications for the control of complex systems. *Cybernetica*, 1(2) , 83-99.

Bar-Yam,Y. (1997). *Dynamics of complex system*s. Massachusetts: Addison-Wesley.

Bennett, F., Koza, J., Keane, M., and Andre, D. (1999). Genetic programming: Biologically inspired computation that exhibits creativity in solving non-trivial problems. In *Proceedings of the symposium on AI and scientific creativity, AISB'99, Edinburgh*, 29 – 38.

Boden, M.A. (1991). *The creative mind: Myths and mechanisms.* London: Weidenfeld and Nicolson.

Bringsjord, S., Bello, P., & Ferrucci, D. (2001). Creativity, the Turing Test, and the (Better) Lovelace Test. *Minds and Machines,* 11 (1), 3–27.

Bringsjord, S. & Ferrucci, D. (2000). *Artificial intelligence and literary creativity: Inside the mind of Brutus, a storytelling machine*, Mahwah, NJ: Erlbaum.

Bringsjord, S., & Schimanski, B. (2003). What is artificial intelligence? Psychometric AI as an answer. In *Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence,* 887–893.

Cohen, H. (1994). The further exploits of AARON, painter. Retrieved November 11, 2008 from http://crca.ucsd.edu/~hcohen/cohenpdf/furtherexploits.pdf

Chomsky, N. (1966). *Cartesian linguistics.* New York: Harper & Row.

Chomsky, N. (1997). Language and problems of knowledge. *Teorema, XVI (2)*, 5-33.

Conant, R. C., & Ashby, W. R. (1970). Every good regulator of a system must be a model of that system. *International Journal of Systems Science*, 1(2), 89-97.

Cosmides, L., & Tooby, J. (2002). Unraveling the enigma of human intelligence: Evolutionary psychology and the multimodular mind. In R. J. Sternberg & J. C. Kaufman (Eds.), *The evolution of intelligence*, 145–198. Mahwah, NJ: Erlbaum.

Dartnall, T. (1994). Introduction: On having a mind of your own. In T. Dartnall (ed.), *Artificial Intelligence and Creativity,* (pp. 29 – 42*),* Dordrecht: Kluwer.

Fauconnier, G. and Turner, M. (2002*). The way we think: conceptual blending and the Mind's hidden complexities.* New York: Basic Books.

French, R. M. (1995). *The subtlety of sameness.* Cambridge: MIT Press.

French, R. M. (1998). When coffee cups are like old elephants or Why representation modules don't make sense. In *Proceedings of the International Conference New Trends in Cognitive Science,* A. Riegler and M. Peschl (eds.), Austrian Society for Cognitive Science, (pp. 158-163).

Gardner, H. (1993). *Multiple intelligences: The theory in practice.* New York: Basic Books.

Godfrey-Smith, P. (2002). Environmental complexity and the evolution of cognition. In R. Sternberg & J. Kaufman (eds.), *The evolution of intelligence, (pp. 233 – 249).* Mahwah, NJ: Erlbaum.

Goertzel, B. & Pennachin, C. (2007). The Novamente *Artificial Intelligence engine.* In B. Goertzel & C. Pennachin, (eds.), (pp. 63 – 129). *Artificial General Intelligence*, Berlin: Springer Verlag.

Gottfredson, L. S. (2003a). Dissecting practical intelligence theory: Its claims and evidence. *Intelligence, 31*(4), 343-397.

Gottfredson, L. S. (2003b). g, jobs, and life. In H. Nyborg (Ed.), *The scientific study of general intelligence: Tribute to Arthur R. Jensen,(* pp. 293-342). New York: Pergamon.

Gottfredson, L. S. (2008). Of what value is intelligence? In A. Prifitera, D. Saklofske, & L. G. Weiss, (Eds.), *WISC-IV applications for clinical assessment and intervention* (2nd ed.), (pp. 545-563). Amsterdam: Elsevier.

Guilford, J.P. (1967). *The nature of human intelligence.* New York: McGraw-Hill.

Harnad, S. (1991). Other bodies, Other minds: A machine incarnation of an old philosophical problem, *Minds and Machines,* 1(1): 43-54.

Hofstadter, D. R. (1985). *Metamagical themas: questing for the essence of mind and pattern.* New York: Basic Books.

Hofstadter, D. R. and Mitchell, M. (1994). The Copycat project: A model of mental fluidity and analogy-making. In K. J. Holyoak and J. A. Barnden, eds., *Advances in Connectionist and Neural Computation Theory: Analogical connections,* 2, (pp. 31–112), NJ: Ablex.

Holland, J. H. (1995). *Hidden order.* Massachusetts: Addison Wesley.

Hutter, M. (2004). *Universal Artificial Intelligence: Sequential Decisions based on Algorithmic Probability.* Springer: Berlin.

Kanazawa, S. (2004). General intelligence as a domain-specific adaptation. *Psychological Review*, 111, 512–523.

Karmiloff-Smith, A. (1992). *Beyond modularity: A developmental perspective on cognitive science.* Cambridge: MIT Press.

Koestler, A. (1964). *The act of creation.* New York: Macmillan.

Koza, J. R., Bennett III, F. H, Andre, D., and Keane, M. A. (1999). *Genetic Programming III: Darwinian invention and problem solving.* San Francisco, CA: Morgan Kaufmann.

Langen, P.H.G. van, Wijngaards, N.J.E. and Brazier, F.M.T. (2004). Designing Creative Artificial Systems, In: A. H. B. Duffy and F. M. T. Brazier (eds.), *AIEDAM, Special Issue on Learning and Creativity in Design,* 18(4), 217-225.

Langley, P., Simon, H. A., Bradshaw, G. L., & Zytkow, J. M. (1987). *Scientific discovery: Computational explorations of the creative processes*. Cambridge: MIT Press.

Legg, S. & Hutter, M. (2007). Universal intelligence: a definition of machine intelligence. *Minds & Machines,* 17(4), 391–444.

Legg, S. & Hutter, M. (2006). A formal measure of machine intelligence. In *Proceedings of the Annual Machine Learning Conference of Belgium and The Netherlands (Benelearn-2006),* Ghent, 73 – 80.

Lenat, D. B. & Guha, R. V. (1990). *Building large Knowledge-based Systems: Representation and Inference in the Cyc project.* Massachusetts: Addison Wesley.

Loosemore, R.P.W. (2007). Complex systems, artificial intelligence and theoretical psychology. In B. Goertzel and P. Wang, (eds,) *Proceedings of the 2006 AGI Workshop.* IOS Press, Amsterdam, 159-173.

McCarthy, J. (2004). What is artificial intelligence? Retrieved November 11, 2008, from www.formal.stanford.edu/jmc/whatisai/whatisai.html

Mithen, S. (1996). *The Prehistory of the mind: A search for the origins of art, religion and science.* London: Thames and Hudson.

Mitchell, M. (2001). Analogy-making as a complex adaptive system. In L. A. Segel and I. R. Cohen, (eds.), *Design Principles for the Immune System and Other Distributed Autonomous Systems*, 335-359, New York: Oxford University Press.

Morowitz, H. & Singer, J. (1995). (eds.). *The Mind, the Brain, and Complex Adaptive Systems.* Massachusetts: Addison Wesley.

Nilsson, N. (2007). The Physical Symbol System Hypothesis: Status and Prospects. In M. Lungarella, et al., (eds.), *50 Years of AI*, Festschrift, LNAI 4850, (pp. 9-17), Springer.

Pinker, S. (1997). *How the mind works.* New York: Norton.

Romero, J., & Machado, P. (2008). (eds.). *The art of artificial evolution.* Berlin: Springer.

Ritchie, G. (2001). Assessing creativity. In Wiggins, G. (ed.), In *Proceedings of the AISB'01 Symposium on Artificial Intelligence and Creativity in Arts and Science*, 3–11.

Ritchie, G. (2007). Some empirical criteria for attributing creativity to a computer program. *Minds and Machines,* 17 (1), 67-99.

Schweizer, P. (1998). The truly total Turing test. *Minds and Machines*, 8(2), 263–272.

Simonton, D. K. (1994). Greatness: *Who makes history and why.* New York: Guilford Press.

Sternberg, R.J. & Lubart, T. (1996). Investing in creativity. *American Psychologist*, 51(7), 677-688.

Sternberg, R.J. & Grigorenko, E.L. (2001). *Practical intelligence and the principal.* Publication Series No. 2, Yale University. Retrieved November 30, 2008 from http://www.temple.edu/lss/pdf/publications/pubs2001-2.pdf.

Sternberg, Robert J. (2005). The Theory of Successful Intelligence. *Interamerican Journal of Psychology*, 39 (2), 189-202.

Turing, A.M. (1950). Computing machinery and intelligence. *Mind*, 59, 433-460.

Wechsler, D. (1944). *The measurement of adult intelligence (3rd ed.).* Baltimore: Williams & Wilkins.