Artificial Intelligence and *Götterdämerung*: The Evolutionary Paradigm of the Future

Kenneth E. Kendall
Rutgers University

Abstract

The impractical metaphor that boldly proclaims information technology is the "information highway" is unsustainable, since information technology is evolving faster than intraplanetary transportation technology. While many of our generation embrace this metaphor and believe that the future is embedded in a network paradigm (web, neural, or other), many of us yearn for the paradigm just beyond it. This paper attempts to predict and define what future paradigms are in store for artificial intelligence (AI). Although we mortals have very few discernible (and replicable) methodologies for predicting the future, we tend to develop our hopes for the future around an identifiable hero. Stories surrounding this hero make up a mythology, so consequently I chose mythology as my methodology.

In particular I depend on Wagner's opera tetralogy referred to Der Ring Des Nibelungen (The Ring of the Nibelung), *which tells the past and prologue of a world full of gods, mortals, giants, dwarves, and (finally) introduces a genetically-evolved hero called Siegfried. In this paper, I make the case that Siegfried is the autonomous agent we refer to in the world of artificial intelligence. Arguably, by observing Siegfried's actions and understanding the symbolism and metaphors in the third opera in The Ring, we can comprehend the future characteristics of autonomous agents. The final opera,* Götterdämmerung, *reveals how the autonomous agent will undergo change, develop, and act.*

The "evolutionary agents," (which I hope to demonstrate is a more appropriate name than "autonomous agents") will: 1) be allowed (and encouraged) to act without representations; 2) be capable of building theories and creating a world of their own; 3) be able to assume any virtual identity they wish; 4) possess free will; and 5) be capable of developing a moral code and a value system of their own. However, the future prospects for evolutionary agents are not very bright, because emotion and love cannot be programmed for it. The final opera of The Ring is called Götterdämmerung *which translates to* The Twilight of the Gods *or, in our world of digital computers,* The Twilight of the Computer Gurus.

Keywords: artificial intelligence, autonomous agent, cognition, evolution, evolutionary agent, free will, *Götterdämerung*, guru, information systems, information technology, mythology, neural nets, opera, virtual reality

ACM Categories: H.5.1, I.2, I.2.0, I.2.3, I.2.4, I.2.6, J.5, K.2, K.4, K.4.2, K.7

Introduction

My great-grandfather was able to foresee the future. Or so say the tales recounted to me by my family down through the generations. Having been trained in the rational school, however, I postulate that perhaps he was the only person in his village who acquired and read books on technology. Now here I stand, less than five years before the new millennium, in the position to predict what will happen in information systems and information technology fifty years hence. I assume that you are well read (in the broadest sense of course, so films, videos, and Internet surfing count, as does the reading of traditional print). Therefore, I recognize that I can't rely solely on my personal ability to keep slightly ahead of the latest technological developments. Perhaps I must truly try to imagine and describe as H.G. Wells was able to do, in 1895.

Methodology

In 1991, I co-authored a paper titled "SEER: A Divergent Methodology Applied to Forecasting the Future Roles of the Systems Analyst, Elaboration, and Review," (Kendall, J. E., et al., 1992). SEER was developed as a method diametrically opposed to the principle of the Delphi technique. With Delphi, experts are queried about a particular problem. These experts answered anonymously, and their answers are shared with the group. After a number of iterations, the Delphi process continues until a consensus is reached. SEER reverses the process, by working backward from the solution and defending our ideas face-to-face until there is sufficient separation of ideas. In the case I described, the result was four distinct scenarios.

Four roles were identified for the systems analyst. Scenario One (Steve Smithson) illustrated that the analyst has become overburdened, being asked to serve in diverse roles such as troubleshooter, change agent, anthropologist, lawyer, and librarian, and suggested the coming of the "renaissance man" who could take on a meta-role of analysts.

In Scenario Two, I became concerned about the overbuilding of systems for their own sake and suggested the analyst take on a role of appraiser and salvager of the information resource and noted that "the appraiser and salvager will needed to decide which entities will survive."

Scenario Three (Ian Angell) suggested that the analyst should reject the emphasis on "positive usability" and emerge with "the ability to deny the virtue of objectivity in all but the most stable environments, and a willingness to place more confidence in human judgment,"(p. 130).

Finally, Scenario Four (Julie Kendall) foresaw a paradigm shift in IS research towards "examining the interactions among people and how they are supported by technology within organizations" and described the analyst as a deep-sea explorer in a fluid organization. In the end we decided "The individual scenarios represent differences in emphasis, rather than opposing views," (p. 133).

Five years after developing it, I still believe that SEER remains a useful methodology for predicting 10-20 years into the future, and that it is useful in preparing for a variety of scenarios the future may bring. However, deduction (or cognition) alone is not powerful or prescient enough to serve as a methodology for foreseeing the future.

In 1964 Buckminster Fuller and his colleagues successfully captivated many urban leaders and dwellers by proposing "The Harlem River Project," in which they envisioned fifteen "colossal beehives" each to be a hundred stories, to spring up from the embattled Harlem urban landscape (Corn and Horrigan, 1996, pp. 54-55). When you see the beehive drawings you will immediately notice that they look distressingly similar to nuclear reactors such as those at Three-Mile Island. Although widely embraced, nothing has been done on these projects; and as Corn and Horrigan point out, it is probably just as well since they are highly reminiscent of the huge urban projects such as

Pruitt-Igo and others which we have finally had to demolish, admitting their failures to address urban problems and their seeming propensity to increase them.

If you would like to see another example of deductive reasoning for predicting the future without a model, examine what Dvorak (1994) predicts about the future of the computer industry and you will realize that too many of the deductive predictions he made have not come true.

In predicting the future of information systems and technology 50 years hence, I believe an overarching methodology is more beneficial. What was needed was a methodology that could be analytical and, at the same time, qualitative – one that could be recognized as descriptive, predictive, and even normative by others. Finally, it had to be enduring. I needed a methodology that would be accepted by many, encourage some, haunt others – one that would be embraced by scientists, poets, philosophers and ordinary people. For these reasons I turned to mythology for my methodology.[1]

I believe that the best predictions of the future are found in ancient stories and myths. My beliefs resonate with those of the historical writer and critic Warner (1995) who in expressing her belief in the importance of myth, quotes the last ideas of Plato's Republic: "And thus it is, Glaucon, that the myth has been saved from oblivion and has not been lost. If we put our faith in it, it may even save us, ourselves."

Then the question turned to which mythology was appropriate? Was it Greek, Roman, Egyptian, Babylonian, Indian, or Arthurian? According to Joseph Campbell (1968):

> "Mythology has been interpreted by the modern intellect as a primitive, fumbling effort to explain the world of nature (Frazer); as a production of poetical fantasy from prehistoric times, misunderstood by succeeding ages (Müller); as a repository of allegorical instruction, to shape the individual to his group (Durkheim); as a group dream, symptomatic of archetypal urges within the depths of human psyche (Jung); as the traditional vehicle of man's profoundest metaphysical insights (Coomraswamy); and as God's revelation to His children (the Church). Mythology is all of these. The various judgments are determined by the viewpoints of the judge. For when scrutinized in terms not of what it is but of how it functions, of how it has served mankind in the past, of how it may serve today, mythology shows itself to be as amenable as life itself to the obsessions and requirements of the individual, the race, the age" (p. 382).

Looking for a rich source that went beyond narrative, I chose to base my predictions on Norse mythology as interpreted by the composer Richard Wagner in his opera tetralogy *Der Ring Des Nibelungen* (*The Ring of the Nibelung*).[2]

In the end, many methodologies designed to help us understand ourselves and illuminate our world fall disappointingly short. Using myths to reveal human unconscious was an important part of Wagner's thinking. Lee (1994, p. 43) notes that the great composer was wary of writing detailed explications of the intended meanings of his operas, believing that myths themselves were the most powerful way to bring true understanding to the human unconscious.

The Ring is based on Wagner's own creative extensions and imaginings from Norse mythology and other stories. It is about power, greed, love, death, heroes, and the transformation of the world.[3] It dramatizes the universe as we

[1] We all use folklore and mythology to explain how computer systems function. Evidence of this is found in the *Windows for Dummies* series (substitute UNIX, the Internet, HTML, Quicken, or whatever, for "Windows" in the title) and user group discussions on the Net. I have proposed and implemented a system called FOLKLORE that documents the important artifacts, symbols, sayings and tales about computer programs in a structured manner (Kendall and Losee, 1986).

[2] This paper concerns Wagner's operas, not Wagner, the man. Forman notes, "There are a few I-can't-stand-Wagners still around as these are usually people who confuse Wagner the nasty little man – and he certainly was racist, chauvinistic, a liar, a megalomaniac with fascist tendencies – with Wagner the composer of genius.," (1994, p. 566).

[3] There are many social, psychological, political, professional, and personal interpretations of The Ring. George Bernard Shaw saw The Ring as a socialistic statement against capitalism, others have used it as a Fascist manifesto (as described in Forman, 1994, p. 556). I have made no political statement this article.

know it from the beginning to the end. (See Lascelles, 1987, and his edited book *Kobbés Complete Opera Book* for a thorough grounding in The *Ring*. All of my spellings of the *Ring Cycle* operas are based on this seminal tome.)

The brilliance of The Ring occurs not just in the wonderfully exquisite music such as the *Ride of the Valkyries* (now familiar to most of the world because of its countless use in films and television commercials), or in the poetry of the libretto, written entirely by Wagner himself. It also resides in the symbolism and metaphors of the drama.

In the information systems world, symbols and metaphors are forever prevalent. My interest in symbolism and information systems is not ephemeral since for many years I have tried to interpret commonly used symbols and predominant metaphors in organizations. My earlier co-authored work includes the interpretation of artifacts and symbols in the decision maker's environment (Kendall and Kendall, 1981) and their impact on decision-making style (Kendall and Kendall, 1984). Later, we examined the use of metaphors to help the systems analyst match method with metaphor to analyze and describe a system better (Kendall and Kendall, 1993) and to determine the success of a proposed system (Kendall & Kendall, 1994).

If decision makers, organizations, and industries rely on mythology (both common, heroic mythology, and organizational myths and legends) to heuristically change the course of their actions, we will be hindered if we observe only the organization's or industry's main characters and their behavior. We must also examine, analyze, and try to interpret the ongoing symbols and metaphors.

Fortunately, the corporate world provides us with abundant symbols and diverse metaphors, and Wagner has reoccurring symbols and metaphors in his music (these are called leitmotifs) and countless metaphors and symbols in his libretto. Are these organizational myths or common myths dynamic or stable? Donington asserts that, "Myth is true for all time. Symbols explain themselves," (1990, p. 101). Throughout this paper, I will try to interpret Wagner's magnificent symbols and archetypal metaphors in order to draw an analogy with the past, present, and future of artificial intelligence. Using The Ring as an analogy for the future of information systems and information technology is ultimately useful because (1) few of you remember your Norse mythology, (2) if you accept the assumption that The Ring foretells the future, then you will probably concur with my predictions, and 3) few of you will attempt to replicate my analysis because anyone who seethes with impatience while waiting a nanosecond to download a file from the Net would be reluctant to endure 17 hours of The Ring to dispute a point!

Analysis

We humans sometimes think we are the center of the universe, not only in space, but also in time. We feel that the principles of management were discovered in our lifetime, that the development of technology is unique to this century, and that information is a powerful resource harnessed with the introduction of the digital computer. We often fail to realize that these elements can be found throughout history, although they have not always been in the foreground. Derry and Williams (1993) point out that technological achievements of early civilizations were not really brought about by man's fear of the divine, but instead by the natural appetite for power. The Ring clearly expresses the belief that greed and power have changed our world.[4]

In The Ring, Wagner describes humans who are common mortals and have little real relevance, but also describes three other orders of players, outside of time – giants, dwarves, and gods. The highest order, however, is reserved for a hero, born of brother and sister, half god-half mortal (Martin, 1984). This character, who is pure in heart, is Siegfried (who doesn't make it onto the stage until the third part of the story). Siegfried is, therefore, our future.

Table 1 interprets the characteristics of each of the primary players in The Ring and maps them to the present computer world as we know it and extends them to a new artificial intelligence paradigm that I believe will be our future 50 years from now. The characters are classic. Even Bill Gates (1995) in *The Road Ahead* boldly entitles

[4] For a revealing glimpse of symbols representing good and evil in the corporate world see J. Kendall (1993).

Characters in *The Ring of the Nibelung*	An interpretation of each Ring character's trait or behavior	Who's who in the computer world-past, present, and future
Dwarves	Capitalists who want power, use money to get it	Information capitalists; high tech firms
Giants	Conservative, non-intellectuals who want to live comfortably	The conservative corporate world; builders of infrastructure
Gods	Intellectuals, who want to be fair and just, who want order in a chaotic world	The gurus of the computer world
Humans	Common mortals who are decent and honorable but break rules set by the gods	End users
Siegfried	A fearless hero, uncorrupted by society	An autonomous agent; (later becomes the evolutionary agent)
Brünnhilde	Siegfried's better half (without her there is no salvation)	Emotion and love, not found in the autonomous or evolutionary agent

Table 1. Comparing the World of *The Ring of the Nibelung* with the Computer World.

chapters "Friction-Free Capitalism" and "Race for the Gold." (*N.B.* In his stage directions for *Das Rheingold*, Wagner describes Wotan, the chief of the gods, as an impetuous, red-haired youth. Any similarity to the CEO of Microsoft is strictly coincidental.) Now that you have met the cast of characters and have glimpsed a map of where we are going, let's draw three quick sketches presenting three currently accepted artificial intelligence paradigms: the Cognitivist Paradigm, the Connectionist Paradigm, and the Action-Selection paradigm. (I am indebted to Franklin, [1995] and Varela et al. [1991] for their excellent books from which these three paradigms are quoted.)

In order to interpret the meaning of the paradigms for the future of information systems, information technology, and society, I first examine the Cognitivist Paradigm, as developed in the myths of *Das Rheingold*; the Connectionist Paradigm, as found mythically in *Die Walküre*; and the Action-Selection Paradigm, as it appears in the myths of *Siegfried*.

I believe that a new, fourth paradigm will inevitably emerge. I have named this emerging paradigm the Evolutionary Paradigm, and find that it departs markedly from the other foregoing artificial intelligence paradigms. The final opera in Wagner's Ring cycle, *Götterdämmerung*, reveals the death of the old self and the birth of a new one; the death of an old world, accompanied by the birth of a new one. If development of the new, fourth paradigm follows the vision set forth in *Götterdämmerung*, then this will be our future.

The Cognitivist Paradigm

A straight forward and streamlined way to compare among the three paradigms, the basic tenets concerning cognition, and its proper functioning, as well as an apt way to introduce the new, fourth paradigm I am proposing, is to quote the efficiently succinct question and answer sections as originally written by Varela et al. (1991) and also used by Franklin (1995):

Question 1: What is cognition?

Answer: Information processing as symbolic computation – rule-based manipulation of symbols.

Question 2: How does it work?

Answer: Through any device that can support and manipulate discrete functional elements – the symbols. The system interacts only with the form of the symbols (their physical attributes), not their meaning.

Question 3: How do I know when a cognitive system is functioning adequately?

Answer: When the symbols appropriately represent some aspect of the real world, and the information processing leads to a successful solution of the problem given to the system.

Applications in the cognitive realm include a wide range of techniques beginning with expert systems and progressing to more complex knowledge-based systems, all of which are based upon symbolic techniques. From the early expert systems programmed in LISP, such as MYCIN and DENDRAL, to those developed using commercial expert systems shells, such as EXSYS and Insight, the main thrust has been to capture the knowledge of the expert using rule-based logic. Classic sources covering applications progress through the middle to late 1980s include Rauch-Hindin (1986) and Martin and Oxman (1988).

For a well-rounded exposition of early knowledge-based systems applications including investment planning, loan approval, auditing, customer service, computer configuration, insurance claim evaluation, inventory planning, media strategy selection, and planning strategic company objectives, see the comprehensive text by Mockler (1989). For a discussion of a failed (and reworked) expert system that typifies approaches of that era, read the fascinating tale told by Markus and Keil (1994). Realizing the limitations of the cognitivist paradigm, many researchers turned to working with neural networks.

The Cognitivist Paradigm and *Das Rheingold*

The three orders, dwarves, giants, and gods (who compose the highest order), deny love and thus set evil into the world. The dwarf seizes gold from the Rhine maidens. This gold, when made into a ring, gives the possessor power over all the world. Wotan, ruler of the gods, steals the ring to pay for Valhalla, his castle built by two giants. One giant kills the other and then turns himself into a dragon (as will be apparent in the fourth opera).

Notice there are no humans in this first *Ring* opera. The technology industrialists (dwarves) seize data (the gold), transform it into information (the ring), and gain power over the world. CEOs (Giants) build the infrastructure (Valhalla) in order to reap the benefits of wealth, that is, to have a comfortable life. Businesses compete, merge, and destroy one another and then turn as conservative and slow moving as a large, sluggish dragon. Computer gurus (the gods), capture the power of information, at least for now.

The gurus have the power. They build computer systems including those with AI (artificial intelligence) based on the manipulation of symbols. They create rules of behavior for end users to follow. Expert systems and early attempts at artificial intelligence were deductive and rule based (Newell, Shaw, and Simon, 1957). That was the past.

The Connectionist Paradigm

When reading the tenets of the second paradigm (quoted directly from Varela et al., 1991), notice that the answers to the standard questions are different, since in this scenario the relationship of simple components is what is creating greater intelligence:

Question 1: What is cognition?

Answer: The emergence of global states in a network of simple components.

Question 2: How does it work?

Answer: Through local rules for individual operation and rules for changes in the connectivity among the elements.

Question 3: How do I know when a cognitive system is functioning adequately?

Answer: When the emergent properties (and resulting structure) can be seen to correspond to a specific cognitive capacity – a successful solution to a required task.

Real world applications of neural networks seem to be paying off almost immediately. One useful

application of neural nets by the bank card industry assists in verifying whether a particular purchase is authentic. The neural network systems used by both MasterCard and Visa International used statistical models built to evaluate details of purchases such as where the item is purchased and what kind of purchase it is. If the transaction is flagged, a bank that holds the card can intervene at that point by phoning the card holder to verify that the cardholder has indeed purchased the item flagged (Murphy, 1996). In attempting to cut down on credit card fraud before it begins, J. C. Penney is creating neural networks to screen people applying for credit cards.

Other types of cheaters are also being found out via neural nets. In Canada, data from historical cases of noncompliance with Canadian tax return laws are being used to train neural networks to assist with predicting the possibility of tax cheating in future cases (Delahanty, 1996). Additionally, neural networks have been used for target marketing, credit card application ratings, forecasting financial climates, assessing risks for insurance underwriting, and also to predict bankruptcy (Foltin and Garceau, 1996). Neural networks are also being used for data mining to uncover "hidden layers of logic" in data that have been warehoused by a particular company (Ross, 1996).

Other business applications now include one called ModelMax, which performs several network modeling tasks such as target mailings. It can also burrow through an enormous number of customer attributes and learn the relationship between variables such as housing type and purchase decisions, together with looking at measures like whether there is a buy or no buy decision and so on (de Ville, 1996).

The Connectionist Paradigm and *Die Walküre*

In *Die Walküre* (*Valkyrie*) two mortals, Siegmund and Sieglinde (twin brother and sister) flee together in an incestual relationship. Wotan commands one of his daughters, Brünnhilde, not to aid them, but she disobeys and saves the pregnant Sieglinde. Wotan puts Brünnhilde to sleep for disobeying him but surrounds her by a wall of fire for her protection. Only a brave man can come through the fire to save her.

In the 1980s and 1990s, end users (common mortals) felt empowered to develop their own information systems (against the laws the gods set forth). Some gurus aided end users, of course, by developing personal productivity tools suitable for use on the personal computer and the Internet. (An interesting chronicle of how this was done from the Microsoft perspective is given in Gates, Myhrvold, and Rinearson, 1995).

The concept of decision support systems (DSS) encouraged end users to develop their own systems. A DSS could handle multiple conflicting goals, improve the decision making, and reduce complexity (Kendall and Kendall, 1995) Complexity, in fact, has become a major concern. For those interested in some lively discussions about chaos and complexity there are many excellent starting points (see Coveney and Highfield, 1995; Gleick, 1987; and Waldrop, 1992).

The preferred approach to developing artificial intelligence in the age of the connectivists is obviously neural networks. Neural nets are based (roughly) on the nervous system, but also appear as a network of end users working towards the main goal of an organization. Neural nets are not programmed in the same way as deductive systems in the Cognitivist paradigm. Rather, they are trained.

End user development and neural nets both are incestual relationships. Recently, I passed around a disk containing spreadsheets and a database to a group of CIOs. I asked them what they feared the most when handing the disk from person to person; and of course, a "virus" was the universal answer. After I told them it was virus-free, no one volunteered another answer. Yet the disk I passed around had spreadsheets that contained faulty assumptions and incorrect equations, while the database contained old and inaccurate data. Would they have verified the assumptions, equations, and data? The point here is that end user development and neural nets contain many problems due to their composition; which is chiefly local rules and connectivity.

When end user systems became susceptible to dangers from intrusion, they asked for the gurus to help. Gurus set up "firewalls," as many prefer to call them, in order to protect end users from

other, unauthorized end users. Users of neural nets will soon ask for the same firewalls. Stoll (1996, p. 3) even proclaims that "They (computer networks) isolate us from one another and cheapen the meaning of actual experience." This is where we are presently.

The Action-Selection Paradigm

According to Franklin (1995), a new paradigm is emerging. He considers this paradigm to be an extension of the enactive paradigm by Varela et al. and names it the action-selection paradigm. He describes it in the following manner:

Question 1: What is cognition?
Answer: The process by which an autonomous agent selects actions.
Question 2: How does it work?
Answer: Actions emerge from the interaction of multiple, diverse, relatively independent modules.
Question 3: How do I know when a cognitive system is functioning adequately?
Answer: When it successfully satisfies its needs within its environment.

One of the most frequently mentioned applications for autonomous agents is that of creating smart interfaces that obligingly adapt to individual preferences for working on the World Wide Web. One Java-based toolset for software developers on the server side that is currently being beta tested is called SDK-100 Srx's Developer's Kit by Selectica Inc. (Bowen, 1996).

Bowen gives the example of a network manager who would use an autonomous agent developed with these tools to get into a vendor's site and accomplish something very practical. The network manager could explore the costs of developing a handful of differently configured networks, using different numbers of hubs, switches, routers and such. The special user interface created via this toolset would permit the user to get graphical displays of all interesting options that fit specifications. Once finalized, the purchase contemplated here could be entered into a relational database, thus cutting out costly delays or possible inaccuracies by capturing the information as close to the sources as possible.

The possibility of developing lifelike autonomous agents has also been explored in the entertainment industry (Maes, 1995). For many readers who have studied the adoption of new technologies, the early appearance of an intelligent agent as an entertaining diversion does not come as a surprise. Virtual environments that permit full-body wireless interaction between a human being and an animated agent are the core of a project called ALIVE which stands for the Artificial Life Interactive Video Environment.

The Third AI Paradigm as Described in *Siegfried*

The third Ring opera is named after its hero Siegfried. Siegfried, son of Siegmund and Sieglinde (the two humans who died in the preceding opera) is the new hope for the world. He is described as the one who knows no fear. Siegfried was raised by a dwarf (after all, dwarves have good qualities too) and he is not very smart in the beginning. We find Siegfried taming and bringing home bears, until he achieves whatever goal he is pursuing, and then he backs the bear offstage telling it "I don't need you anymore." He then forges an indestructible sword and calls the sword "Needed." He uses it to kill a dragon (which is really a giant disguised in dragon's clothing), talks to the birds, captures the ring, walks through the fire, and awakens Brünnhilde.

Siegfried is the autonomous agent proposed by Brustoloni (1991). Siegfried is capable of autonomous, purposeful action that responds to external stimuli. Furthermore, Seigfried is an adaptive agent, which makes him capable of acquiring domain knowledge. Seigfried accomplishes many goals, but in doing so he starts out slowly (knowing nothing about his real parents), forges a sword (called "Needed") which allows him to takes other actions like slaying a dragon. Adaptive agents are capable of performing actions they were not previously capable of performing by acquiring knowledge and building on earlier actions.

Siegfried is not an adaptive planning agent (he seems dim at times) but instead is as Brustonlini describes; an adaptive case-based agent. In The Ring, Seigfried exhibits all of Brustolini's behaviors: instinctive, habitual, problem-solving, playing, and theory making as he tames a bear, forges a sword, slays a dragon, talks to the birds, and walks through the ring of fire to awaken Brunnhilde.

Seigfried is also the autonomous agent described by Franklin (1995) in his action-selection paradigm. It appears to me that Siegfried also exhibits the six behavioral characteristics of Newell and Simon (1972, pp. 806-890):

1. *Interruptibility*. Siegfried would return to the action after being distracted (usually by the dwarf).
2. *Subgoaling*. Siegfried could be interrupted to work on a sub-goal and then return to his main goal, to rescue Brünnhilde.
3. *Depth-first subgoaling*. The path to Brünnhilde had many sub-goals. Siegfried had to kill a dragon, which implies he had to forge an indestructible sword (he himself called the sword "Needed," which in turn implies he needed to know smithing.)
4. *Equifinality*. Siegfried attempts to talk to a bird. He makes bird sounds, failing that, he then tries blowing on reeds, followed by attempting to blow his horn. Finally he drinks the dragon's blood which allows him to talk to the bird. (This sounds like an adventure game, but the principle of equifinality means that if at first you don't succeed, try, try, again).
5. *Avoidance of repetition*. Siegfried went on to do many different things. Siegfried's use of recall to avoid repetition is important here, since in the fourth opera he loses his memory.
6. *Consummation*. Once Siegfried's goal of rescuing Brünnhilde was realized, all effort toward that goal stopped. (Wagner really plays up the consummation scene between Siegfried and Brünnhilde.)

Now that the three paradigms and their interpretation, as interpreted through the myths present in the first three operas of Wagner's Ring cycle, have been examined, allow me to introduce a new fourth paradigm, the Evolutionary Paradigm, whose meaning can be understood through the myths in the final opera of the *Ring* which is titled *Götterdämmerung* (or *The Twilight of the Gods*).

The Evolutionary Paradigm

The fourth and final paradigm will be the evolutionary paradigm. It is destined to develop as a post-representational paradigm. As we move quickly through a comparison of answers to the standard questions posed to the first three paradigms, take into consideration the shifts that are occurring. I based the questions and answers in this section on the format used by Franklin (1995) and they are:

Question 1: What is cognition?
Answer: The process in which an autonomous agent acts without representations, builds theories, assumes virtual identities, and exhibits free will.
Question 2: How does it work?
Answer: Through evolving systems that can access the merits of inconsistent motives, can store and compare motives, change upon self-evaluation, and adapt to approval or disapproval of external agents.
Question 3: How do I know when a cognitive system is functioning adequately?
Answer: When the autonomous agent evolves and creates a new world with its own moral code and values.

One of the reasons for eliminating representation is that gurus have biases. George Reeke, Jr. and Gerald Edelman (as quoted in Franklin, 1995, p. 365) say:

> "Once an appropriate representation is available, many problems do become amenable to automatic solution. In our view, however, the problem requiring intelligence is the original one of finding a representation. To place this problem in the domain of the systems designer rather than in that of the designed system is to beg the question and reduce intelligence to symbol manipulation."

Consequently, an agent, allowed to evolve and exist in a virtual world without representations, would have no biases built in. This evolutionary agent then would not be one of the computer gurus or system designers, but a product of the gurus some generations removed. (Siegfried, was the son of brother-sister parents whose father was the god Wotan and whose mother was a mortal. Siegfried was raised by a dwarf.) Furthermore, the evolutionary agent must also act of free will.

A fourth paradigm is destined to emerge. Although my reading of Newell and Simon

(1972) implies that they are steadfastly against the representationless paradigm, their position will not prevail. Too many entrepreneurs and (eventually) end users will find the seductive nature of a virtual, representationless world overpowering. I will even go so far as to suggest that it will be analogous to the phenomena described by Nye in *American Technological Sublime* (1996). Nye explores what he refers to as the "electrical sublime," from its introduction to its high point as the central technology theme for the city of Buffalo's Pan-American Exposition of 1901, to the commercialization resulting in undesirable electrical signs, then finally through the consumer sublime of Las Vegas and Disneyland. Where at one time the "technological sublime had exhorted the observer to dominate and control nature . . ." (compare this to the first AI paradigm), the consumer sublime "privileges irrationality, chance, and discontinuity," (compare this to the fourth, representationless AI paradigm).

I believe that autonomous agents in the fourth, or evolutionary, paradigm will:

1. be allowed (and encouraged) to act *sans* representations.
2. be able to create a world of their own, by building theories about their world.
3. be able to assume any virtual identity they wish (or act as an artificial clone of a human).
4. possess and exhibit free will.
5. be capable of developing a moral code and a value system of their own.

But if a representationless world can exist, the autonomous agent in this paradigm will by nature possess certain traits and features. For help in identifying these unique traits, we can turn to a set of design characteristics as posed by Sloman (1988). Franklin (1995) selected ten of the more critical design considerations and these can be found in Table 2, which also points out differences between the third and fourth paradigm. Using the language of Sloman, I conclude that the evolutionary agent in the fourth paradigm will be distinguished from the autonomous agent in the action selection paradigm in the following ways:

1. The evolutionary agent can assess the merits of different inconsistent motives, while the autonomous agent actions are controlled by the most recently generated motive. This, I believe, is a necessary but not sufficient condition for the existence of an evolutionary agent. A representationless world would be inconsistent (and sometimes irrational, discontinuous, and even left up to chance) by us humans.
2. The evolutionary agent can simultaneously store and compare different motives, while the autonomous agent takes on singular motives sequentially. I believe that theories can be built and worlds can be created only when motives are evaluated and compared in parallel rather than in sequence.
3. The evolutionary agent can change under the influence of genetically determined factors, while the autonomous agent can change only in the light of interactions and the inferences drawn therefrom. I would argue that the evolutionary agent has the power to change at will either because of some genetically determined factor (Franklin [1995] suggests that this can mean an event such as puberty, but I think Dawkins [1987], as interviewed by Bass [1994], might interpret this in a very different manner) or by a factor made available, for example technology such as virtual reality. In this way the evolutionary agent can assume a different virtual identity.
4. The evolutionary agent's motive generators and comparators are accessible to explicit internal scrutiny, analysis, and change, while the autonomous agent's motive generators are merely uncontrolled side effects. I believe that this design characteristic is a necessity for free will.
5. The evolutionary agent's motive generators and comparators change under the influence of the likes and dislikes, or approval/disapproval of other agents. Autonomous agents, however, are influenced solely by how things affect it. I conclude that this design characteristic is necessary for the evolutionary agent to be capable of developing a moral code and a value system.

To reiterate, the evolutionary agent must possess certain characteristics that distinguish itself from the autonomous agent as we describe it today. Some of these design characteristics may be thought to be impossible, or at least improbable, by many; but then again, many thought that flight was impossible. (That is to say the mythology of Daedalus and Icarus came true to some degree, at least.)

Design Consid-eration	Alternative ●	Alternative ○	Action selection paradigm	Evolutionary paradigm
1	Can simultaneously store and compare different motives	Has only one motive at a time	○	●
2	Single top-level goal	Independent sources of motivation	○	○
3	Modification of its motive generators in the light of experience	Generators and comparators are fixed for life	●	●
4	Change under the influence of genetically determined factors	Can change only in the light of interactions	○	●
5	Motive generators and comparators accessible to explicit internal scrutiny, analysis, and change	Motive generators and comparators are merely uncontrolled side effects	○	●
6	Motive generators and comparators change under the influence of likes and dislikes, or approval or disapproval of other agents	Influenced by how things affect it	○	●
7	Agents that are able to extend formalisms they use for thinking	Agents who are not capable	●	●
8	Assess the merits of different inconsistent motives	Controlled by the most recently generated motive	○	●
9	Monolithic hierarchical computational architecture	Independent subsystems can generate different goals	○	○
10	Decisions take on some kind of "democratic" voting basis or a numerical summation	Conflicts are resolved on the basis of qualitative rules	○	○

Table 2. Differences Between the Third and Fourth AI Paradigm as Illustrated using some of Sloman's Design Considerations

There are always people who think that future technologies are impossible. H. G. Wells' "World Brain" was thought to be impossible, but he says people who think that future essentially envisioned a form of the World Wide Web we have today. [5]

In *Being Digital*, Negroponte (1995) states: "Computers are not moral; they cannot resolve complex issues like the rights to life and to death." Within the framework of the fourth AI paradigm, the evolutionary agent is indeed capable of generating its own moral code. Development and applications are still years away, however. [6]

In *Paradigms Lost*, Casti (1989) poses the question "How real is the 'real world'?" He summarizes the work of over a dozen scientist-philosophers who he calls romantic realists (such as Bohr, Wheeler, Heisenberg, Everett and Deutsch) who believe that there is no objective reality, and dogwork realists (like Einstein, Bohm, Bell, and Cramer) who believe that a single, observer-independent reality may exist. Casti weighs the evidence and sides with the multiple world interpretation.

Einstein may be correct that the Newtonian universe is the real universe, but does it matter? If enough scientist-philosophers are convinced that there is no objective reality, then the mainstream design elements of AI in the future will include domains of subjective, or virtual reality. This mandates that representations must be absent at the start, and that the evolutionary agent (formerly autonomous) now be given the chance to build representations and theories and eventually develop its own moral code and values.

[5] For an interesting discussion of Wells' World Brain and his tour to raise money for its implementation, see Campbell-Kelly and Aspray (1996, pp. 285-6).

[6] One of my colleagues, Roger Clark of the Rutgers School of Law - Camden, has appeared at the International Court of Justice in the Hague as counsel for Samoa as the Court is expected to rule on the legality of the use of nuclear weapons. After reading his incisive notes I am not sure humans are capable of discussing such an irrational issue. Perhaps we need the evolutionary agent and its newly-developed moral code for such problems.

The Evolutionary Paradigm and *Götterdämmerung* (*Twilight of the Gods*)

Early in *Götterdämmerung* there is a scene where three sisters who are winding the rope of destiny are asked "Do you know what will be?" Some critics of Wagner are quick to point out that this scene is tedious because the sisters take awhile explaining the past and present, and then just as the future is ready to be revealed, the rope breaks. I think Wagner deliberately inserted this scene in order to highlight the importance of representations, that are the past and present. This scene, therefore, sets up the rest of *Götterdämmerung* (and sets up our fourth paradigm as well).

Siegfried gives Brünnhilde the ring and she gives him in return wisdom and strength and explains that he is ready to go on to obtain new goals. Siegfried journeys down the Rhine, makes new human friends, the brother and sister act Gunther and Gutrune. A half-dwarf manages somehow to get Gutrune to drug Siegfried, making Siegfried forget much of what he knows, especially his relationship with Brünnhilde. And this is where things start to go drastically wrong for our hero.

In our world, the technology capitalists will encourage a virtual reality world, especially in AI, because they will want to regain control of the gold. The computer gurus, of course, would prefer not to lose control, but they will not resist the temptation to create an evolutionary agent that exhibits free will, and the gurus will take on the challenge. The corporate giants will not be players in this development, while the end users will be duped into participating in the development of the evolutionary agent in its virtual world.[7]

At this point Siegfried lives in a representationless world. One of his goals is to find (or define) representations. An early representation of Siegfried shows Gutrune as a thing of beauty and he woos her. He doesn't recognize Brünnhilde, but develops a theory about her.

[7] "The element in the giant's symbolism which is perhaps the most important here has to do with their massive stupidity. Giants are always conservative forces, on the side of the past rather than of the future," (Donington, 1974, pp. 110-111).

Siegfried dons a magic helmet (the tarnhelm), which allows him to appear to be any other being, and he chooses to look like Gunther. Admittedly this is strange, but who would an evolutionary agent choose if it had no prior representations, anyway? Disguised as Gunther, he seeks out Brünnhilde so he can give her to (the actual) Gunther as a bride. Along the way he develops a moral code of conduct and values. Although he thinks he is doing some good, he seems to make matters worse.

Siegfried has free will, his destiny no longer controlled by the gods.[8] The gods at this point are indirectly observing the events. Wotan, the chief god, calls for the return of his two ravens named Reason and Memory (sounds to me as if Wagner predicted the main components of a digital computer) who have been watching Siegfried until the end. Eventually Siegfried is slain, and Brünnhilde joins him on his funeral pyre. The fire grows larger and eventually destroys Valhalla and with it comes the death of the gods. The ring is returned to the Rhine. Brünnhilde needs to be reunited with Siegfried in order for there to be redemption.

If the development of the fourth paradigm parallels *Götterdämmerung*, then the fourth paradigm will be developed, and evolutionary agents will be allowed to act with free will, without representations, and build their own theories, take on virtual identities, and develop their own moral codes and values. The results may be tragic, just as they were for Siegfried (the autonomous, turned evolutionary agent) and Brünnhilde (the emotion and love needed to complement the attributes of the evolutionary agent). The gods, or gurus as we like to call them, perish as well and the glorious music in the end gives us hope that our world as we know it may return to its natural (analog?) state and be reborn.[9]

Brünnhilde beta testing the virtual reality helmet of the future.

Drawing by Victoria Roberts, caption by Ken Kendall.
Reprinted with permission from the
Kendall & Kendall cartoon collection, 1995.

Conclusion

I predict that we will someday see an autonomous agent that fulfills the promises of the action selection paradigm.

I offer a bolder prediction: a new, "evolutionary" paradigm will emerge and attempts will be made to introduce an evolutionary agent to a representationless virtual world.

Stoll (1996, p. 148) suggests that "Computer networks replace that lust for the physical with virtual reality," On multimedia displays, he adds "These displays only mock: Grand Canyon Dam, for instance, has a breathtaking photo display of the now submerged canyon. When I saw what the Bureau of Land Reclamation destroyed, I started to cry and had to leave (1996, p.149)."

I can think of countless counterexamples where virtual reality might enable us to be more humane, not perverse as Stoll and others like Slouka (1995) and Turkle (1995) postulate. If a person can experience elephants in the wild, their size, color, odors, texture of their skin, and the sounds they make in a virtual reality world, then perhaps we can deconstruct the concept of a zoo and save animals from a cruel, caged

[8] Lee (1994) observes, "If, then there is one fundamental insight that, intentionally or intuitively, underlies the Ring, it is that everything existent has evolved from one primal substance. And a second important idea is that man, to be come man, had to separate himself from the rest of nature by evolving out of it into a conscious state of being – or as the myths seem always to picture it, by wresting consciousness from nature," p. 41.

[9] Donington (1974) states, "The fantasy of returning to Mother Nature also holds the positive possibilities of going down symbolically into her womb not as a victim of seduction but as a candidate for rebirth," (p. 41).

life.[10] The point here is that the gurus will provide examples and, more importantly, hype that will ensure the development and perfection (undoubtedly!) of virtual reality.

But if it is possible for something (an autonomous or evolutionary agent) to enter a virtual world without memory (or representations) and experience the Grand Canyon after and before the Bureau of Reclamation constructed the dam or experience the very same animals captivated in the once "great" London Zoo and in the wild, what do we have? We have Wells' (1895) concept of a time machine, albeit virtually.

The question that remains is when? When will the evolutionary agent create a virtual world? We know that a rubber band will eventually snap, but it is difficult to predict the exact split second when it will happen. Perhaps Dawkins (1987, p. 133) gives us a clue, "Life doesn't immediately follow from self-replication. Life follows from many generations of gradual evolution after self-replication has begun."

Translating concepts that deal with the theory of hyperspace (Kaku, 1994) reminds us that the most advanced formulation of this is superstring theory and that the most important idea about it is that it will be powerful enough to unify all known physical phenomena in an astonishing simple and elegant framework. One of the hallmarks of the theory is a precise prediction that there are ten dimensions, not three, thus we should no longer be bound by thinking only of three. Kaku also notes that we must wait centuries or millennia to be able to develop capabilities to manipulate space-time, or to reach those who have already tamed hyperspace. He concludes that "any civilization that masters the hyperspace will become lord of the universe," (p. xi).

I caution, however that, even given the comments of Dawkins and Kaku, generations in the computer world are extremely short indeed.

Once this is possible, the evolutionary agent is free to build theories and react to observations from this created world. Memes, the basic unit of cultural transmission, will have to come from within the created world, not from our families and friends as in the real world. The evolutionary agent can then create its own moral code and values.

Dery (1996) contends that someday we will be capable of "'downloading' our selves after having delegated, one by one, all of our mental and physical functions to our machines." Perhaps this will happen, although I am not certain of this. The only clue Wagner left in *Götterdämmerung* is that Siegfried is indestructible except for a small area of his back. (It is his own "Achilles heel" so to speak, borrowing from yet another myth). Perhaps this signifies that Siegfried has evolved into a machine, and is not in a human form after all. (Recall that in the film "2001: A Space Odyssey," HAL had an Achilles heel as well).

One thing I am certain of is that the evolutionary agent will never have an emotional soul. Depending heavily on my methodology (or perhaps "mythodology") based on Wagner's *Götterdämmerung*, I conclude that our mind and physical functions may be downloaded, but our souls will not. Brünnhilde, complete with love and emotion, is Siegfried's soul.

Although a representationless world will be created, the value is not clear to me. Relationships are also symbols. "Not only do symbols represent inner realities; they are inner realities. The power over good or ill can be great: the Cross, for example; or the Swastika," (Donington 1974, p. 21).

The problem is something is bound to go wrong. Brodie (1996) worries about mind viruses and explains, "Once created, a virus of the mind gains a life independent of its creator and evolves quickly to infect as many people as possible." Although this may be possible if an evolutionary agents creates its own world, I do not foresee this happening.

Given the assumption that relationships are symbols, we again turn to Donington (1974, p. 21) who states, "Thoughts are intellectual counters; symbols are emotion counters, com-

[10] I am playing devil's advocate here. In essence Stoll and I come to similar conclusions that virtual reality is lacking emotion and love. Cliff and I both attended astronomy and other classes at the Buffalo Museum of Science and share some of the same representations. We even pronounce the red dwarf star Betelgeuse in the same way.

pounded at once of feeling and of intuition." This moves me to conclude that artificial life without representations are doomed to fail from the very beginning. In essence, the computer guru will not be able to create sustainable life, even though the attempt will be made.

Instead, I see the world misunderstanding the autonomous, or evolutionary agent (as Brünnhilde misunderstood the virtual world Siegfried) and ensuring it is destroyed. The result will be the sacrifice of the evolutionary agent and the twilight of the computer gurus.

References

Bass, T. A. (1994). *Reinventing the Future: Conversations with the World's Leading Scientists.* Reading, MA: Addison-Wesley Publishing Company.

Bowen, T. S. (1996). "Web Searches to Gain Smart User Interface," *Infoworld*, Vol. 18, No. 30, July 22, pp. 1, 24.

Brodie, R. (1996). *Virus of the Mind: The New Science of the Meme.* Seattle, WA: Integral Press.

Brustoloni, J. C. (1991). *Autonomous Agents: Characterization and Requirements.* Carnegie Mellon Technical Report CMU-CS-91-204. Pittsburgh: Carnegie Mellon University.

Campbell, J. (1968). *The Hero with a Thousand Faces*, (2nd Edition). Princeton, NJ: Princeton University Press, p. 382.

Campbell-Kelly, M., and Aspray, W. (1996). *Computer: A History of the Information Machine.* New York: BasicBooks, a Division of Harper Collins Publishers, Inc.

Casti, J. L. (1989). *Paradigms Lost, Tackling the Unanswered Mysteries of Modern Science.* New York: Avon Books, A Division of the Hearst Corporation.

Corn, J. J., and Horrigan, B. (1996). *Yesterday's Tomorrows: Past Vision of the American Future*, Baltimore, MD: The Johns Hopkins University Press, The Johns Hopkins Paperbacks Edition, pp. 54-55.

Coveney, P., and Highfield, R. (1995). *Frontiers of Complexity: The Search for Order in a Chaotic World.* New York: Ballantine Books, a division of Random House.

Dawkins, R. (1987). *The Blind Watchmaker: Why the Evidence Reveals a Universe Without Design.* New York: W. W. Norton & Company, Inc.

Delahanty, P. (1996). "Government Using Neural Nets to Snag Tax Cheaters," *Computing Canada*, Vol. 22, No. 16, August 1, p. 6.

Dery, M. (1996). *Escape Velocity, Cyberculture at the End of the Century.* New York: Grove Press.

Derry, T. K., and Williams, T. I. (1993). *A Short History of Technology: From the Earliest Times to A. D. 1900.* New York: Dover Publications, Inc.

de Ville, B. (1996). "Direct Marketing with ModelMax," *Marketing Research: A Magazine of Management & Applications*, Vol. 8, No.1, Spring, pp. 56-58.

Donington, R. (1974). *Wagner's 'Ring' and Its Symbols,* (3rd Edition). London: Faber and Faber.

Donington, R. (1990). *Opera and Its Symbols: The Unity of Words, Music, & Staging.* New Haven: Yale University Press.

Dvorak, J. C. (1994). *Dvorak Predicts: An Insider's Look at the Computer Industry.* Berkeley, CA: Osborne McGraw-Hill.

Foltin, C., and Garceau, L. (1996). "Beyond Expert Systems: Neural Networks in Accounting," *National Public Accountant*, Vol. 41, No. 6, June, pp. 26-30.

Forman, D., Sir, *A Night at the Opera.* New York: Random House, 1994.

Franklin, S. P. (1995). *Artificial Minds.* Cambridge, MA: A Bradford Book, published by MIT Press.

Gates, W., with Myhrvold, N, and Rinearson, P. (1995). *The Road Ahead.* New York: Viking Penguin.

Gleick, J. (1987). *Chaos: Making a New Science.* New York: Penguin Books.

Kaku, M. (1994). *Hyperspace: A Scientific Odyssey Through Parallel Universes, Time Warps, and the 10th Dimension.* New York: Anchor Books published by Doubleday.

Kendall, J. E., (1993). "Good and Evil in the Chairmen's Boiler Plate: An Analysis of Corporate Visions of the 1970s," *Organization Studies,* Vol. 14, No. 4, pp. 605-626.

Kendall, J. E., and Kendall, K. E. (1993). "Metaphors and Methodologies: Living Beyond the Systems Machine," *MIS Quarterly,* Vol. 17, No. 2, pp. 149-171.

Kendall, J. E., and Kendall, K. E. (1994). "Metaphors and their Meaning for Information

Systems Development," *European Journal of Information Systems*, Vol. 3, No. 1, pp. 37-47.
Kendall, K. E., and Kendall, J. E. (1981). "Observing Organizational Environments: A Systematic Approach for Informational Analysts," *MIS Quarterly,* Vol. 5, No.1, pp. 43-55.
Kendall, K. E., and Kendall, J. E. (1984). "Structured Observation of the Decision-Making Environment: A Validity and Reliability Assessment," *Decision Sciences*, Vol. 15, No. 1, pp. 107-118.
Kendall, K. E., and Kendall, J. E. (1995). *Systems Analysis and Design*, third edition. Englewood Cliffs, NJ: Prentice Hall.
Kendall, K. E., and Losee, R. D. (1986). "Information System FOLKLORE: A New Technique for System Documentation," *Information & Management*, Vol. 10, No. 2, pp. 103-111.
Kendall, J. E., Kendall, K. E., Smithson, S. and Angell, I. O. (1992). "SEER: A Divergent Methodology Applied to Forecasting the Future Roles of the Systems Analyst," *Human Systems Management,* Vol. 11, No. 3, pp. 123-135.
Lascelles, G. (1987). "Der Ring Des Nibelungen (The Ring of the Nibelung), in *Kobbés Complete Opera Book*, (10th Edition) G. Lascelles, The Earl of Harewood (Ed.), London: The Bodley Head, pp. 214-262.
Lee, M. O. (1994). *Wagner's Ring: Turning the Sky Around*. New York: Limelight Editions.
Maes, P. (1995). "Artificial Life Meets Entertainment: Lifelike Autonomous Agents," *Communications of the ACM*, Vol. 38, November, No. 11, pp. 108-114.
Markus, L. M., and Keil, M. (1994). "If We Build It, They Will Come: Designing Information Systems that People Want to Use," *Sloan Management Review*, Vol. 35, Summer, No. 4, pp. 11-25.
Martin, G. (1984). *The Opera Companion.* London: John Murray (publishers) Ltd.
Martin, J. and Oxman, S. (1988). *Building Expert Systems*. Englewood Cliffs, NJ: Prentice Hall.
Mockler, R. J. (1989). *Knowledge-Based Systems for Management Decisions*. Englewood Cliffs, NJ: Prentice Hall.
Murphy, P. A. (1996). "Fraud Hits Store Cards," *Stores*, Vol. 78, No. 1, January, pp. 58-60.
Negroponte, N. P. (1995). *Being Digital.* New York: Vintage Books, a Division of Random House.
Newell, A., and Simon, H. A. (1972). *Human Problem Solving.* Englewood Cliffs, NJ: Prentice-Hall.
Newell, A. J.. Shaw, C, and Simon, H. A. (1957). "Empirical Explorations of the Logic Theory Machine," *Proceedings of the Western Joint Computer Conference*, pp. 218-239.
Nye, D. E. (1996). *American Technological Sublime*. Cambridge, MA: MIT Press.
Rauch-Hindin, W. B. (1986). *Artificial Intelligence in Business, Science, and Industry, Vol. 1: Fundamentals*. Englewood Cliffs, NJ: Prentice-Hall.
Ross, J. R. (1996). "Data Mining: Digging Deeper for Information Treasures," *Stores,* Vol. 78, No. 5, May, pp. 66-68.
Sloman, A. (1995). "Disposing of the Free Will Issue." Connectionist List message on the Internet, June 19, 1988, as quoted in Franklin, S. P., *Artificial Minds.* Cambridge, MA: A Bradford Book, published by MIT Press.
Slouka, M. (1995). *War of the Worlds: Cyberspace and the High-Tech Assault on Reality.* New York: BasicBooks, pp. 55-56, and 153.
Stoll, C. (1996). *Silicon Snake Oil.* New York: An Anchor Book published by Doubleday.
Turkle, S. (1995). *Life on the Screen: Identity in the Age of the Internet.* New York: Simon & Schuster.
Varela, F. J., Thompson, E., and Rosch, E. (1991). *The Embodied Mind.* Cambridge, MA: MIT Press.
Waldrop, M. M. (1992). *Complexity: The Emerging Science at the Edge of Order and Chaos.* New York: A Touchstone Book.
Warner, M. (1995). *Six Myths of Our Time*, New York: Vintage Books.
Wells, H. G. (1996). *The Time Machine.* London: William Heinemann, 1985 as reprinted in *The H. G. Wells Reader*, Philadelphia: The Running Press.

About the Author

Kenneth E. Kendall, a professor of MIS at Rutgers University, School of Business-Camden, has always been fascinated with both mythology and technology. Ken was exposed to mythology at birth by his mother's obstetrician, who was also a scholar of mythology. (Twenty years on, he worked with the good doctor designing systems for regional blood banking). Ken, who loves any new technology, was once denied warranty service on his Apple computer because

it had too few digits in the serial number. He enjoys his work as the MIS Editor for the *Journal of Management Systems*; as a functional editor of MIS for *Interfaces*; as an Associate Editor for *Decision Sciences*, the *Information Systems Journal*, and *Information Resources Management*, and he also welcomes his current role as a vice-president for the Decision Sciences Institute.