

Automating Human Decisions

The Business Case for Common Sense Reasoning™

Written By
Marc Schneiderman

Copyright © 2008
Mobile Agent Technologies

Table of Contents

TABLE OF CONTENTS	2
INTRODUCTION	3
WORLD FULL OF UNCERTAINTY	3
CHANGING BUSINESS MODELS	3
HOW PEOPLE MAKE DECISIONS	4
THE BEAUTY OF THE HUMAN MIND	5
LIMITATIONS TO HUMAN COGNITIVE ABILITIES	5
AUTOMATED DECISION MAKING SYSTEMS.....	6
DRIVING FORCES	7
<i>Business Survey Results.....</i>	<i>7</i>
<i>Core Applications No Longer Provide a Competitive Advantage.....</i>	<i>7</i>
<i>Gap in Traditional Business Intelligence Applications.....</i>	<i>7</i>
UNDERLYING TECHNOLOGIES.....	8
BENEFITS OF DECISION AUTOMATION	8
WHY EXPERT SYSTEMS FAILED.....	9
FIRST GENERATION SYSTEMS.....	10
<i>Current Implementations.....</i>	<i>10</i>
<i>Limitations of Existing Systems.....</i>	<i>11</i>
CONCLUSION	13

Introduction

Except perhaps for our transition from an agrarian society to that of an industrialized one, mankind has not seen the advent of so many new and diverse technologies. These inventions, including the computer and internet, have had a very profound and lasting affect on each and every one of our lives. They have affected the way we meet, collaborate, work, and vacation. As a result of innovative software tools that have been introduced over the past ten years, including business rules engines and data mining toolkits, organizations are now able to make faster, more accurate, consistent, and unbiased decisions automatically.

World Full of Uncertainty

We live in an ever changing world, fraught with uncertain outcomes, where established norms no longer hold true, both from a political perspective as well as in business. For all intents, it is a world turned upside down. Recently Bear Stearns, an eighty-five year old Wall Street investment bank and an underwriter of many successful initial public offerings, had to resort to short term borrowing from the Federal Reserve, and a subsequent “fire sale” to a larger bank in order to avoid insolvency. On the other hand, China, a country under the political rule of a communist government, having state supported industries, is home to the world’s fastest growing economy, offering investment vehicles in some of the most attractive small cap companies.

Changing Business Models

As Pulitzer prize winning author Tom Friedman eloquently stated, the world is flat. Today reservations at JetBlue Airlines are “homesourced” to employees working from their homes in Salt Lake City. The operations of unmanned predator drones on the battlefield in Afghanistan are controlled by air force personnel at Nellis air force base near Las Vegas, Nevada. Information technology outsourcers, such as Wipro, Tata, and Infosys, now design, develop, and maintain, core applications offshore for Fortune 500 companies for less money than would be possible using American workers.

Li & Fung, initially founded as a trading company in 1906, now produces eight billion dollars worth of garments and consumer goods a year, for some of the worlds leading brands. Yet they own no factory, or means of distribution, but rather have mastered the complex art of orchestrating the operations of thousands of different suppliers, each manufacturing or transporting raw materials, fabrics, or finished goods. In actuality, they are not in the clothing business but rather masters at the art of supply chain management and logistics. Indirectly the company provides work for more than two million people across the globe, of which less than one half of one percent are actually on their payroll.

With the advent of the modern day computer, as well as the internet, the age of information has moved into full swing. Newly formed companies like Yahoo, eBay, Amazon.com and Google have all created innovative business models and surpassed their traditional “brick and mortar” rivals. The age of industry brought about a transition from goods being made by hand, to products being manufactured in a factory by machine. The information age, which initially brought about the automation of rote manual tasks, such as typing a letter, or posting a journal entry to the general ledger, is now poised to automate the very nature of human thought. The traditional call center, which was outsourced in 1995 due to cheaper overseas labor rates, will soon be replaced by automated decision making systems located in a data center in Des Moines, Iowa.

How People Make Decisions

Theory on rational thought processes can be traced all the way back to the ancient Greek philosophers Plato and Aristotle. Since the age of discovery, the basic premise of the rational decision making model was that a human will make a decision by first defining a problem, generating all possible solutions, creating objective assessment criteria for all of the available options, methodically comparing all of the choices, and finally choosing the best one.

Benjamin Franklyn, in his autobiography, summed up this approach by saying that he would make decisions by first drawing a “T” on a plain piece of white paper. On the top of one side he would write the word “For”, and on the other side, the word “Against”. He would then list all of the reasons for taking a given action, and all rational for not going forward with it as well. Each of the reasons he would assign weights to as well. At the end he would add up all of the reasons for, and those against, and the course of action to follow was made for him.

Research performed over the course of the past forty years, by noted cognitive psychologists, have proven that people do not actually use the rational decision making model. The reasons behind this conclusion lie with the underlying assumptions that this decision making model is based upon. First of all, it is predicated on the basis that there is a single best solution. Secondly, it requires complete information to be known, including environmental factors, alternatives, goals, and consequences. Next, it is assumed that the execution of the decision making process is linear in nature. Lastly, it takes for granted that there is an unlimited amount of time in which to reach a conclusion.

What is interesting to note though, is that completely inexperienced people in any given field will attempt to utilize this method to make their decisions. While those with many years of experience performing a given task will use an alternative approach known as the bounded rational decision making model. The bounded model assumes that we live in an imperfect world with time constraints, and oftentimes incomplete information available to us. Secondly, there is the notion that problems may be ill defined. Not all of the alternatives can be generated or will be known. Instead of finding the optimal

solution, which may not even exist, the methodology settles on finding the first “workable” solution. It is a non-linear sequence of events, whereby taking a specific course of action may itself alter the very nature of the problem, causing another iteration of the process to take place in a recursive fashion, creating new goals and actions. In Gary Klein’s landmark book entitled “Sources of Power: How People Make Decisions”, he diagrams the thought processes that are involved and calls it the recognition primed decision model.

The Beauty of the Human Mind

Nature has endowed humans a number of incredible mental capabilities that help us to make decisions. First, we have been given intuition, sometimes referred to as gut feelings or common sense, which allow us to make a key decision in a very short period of time, usually within a number of seconds. When asked about how we reach a given decision, such as coming in out of the rain, most people respond that the answer was ingrained into their psychic since childhood, meaning they did not even have to think about it. In reality, what we call common sense is actually many different, albeit rather simple, rules of thumb that we acquire over the course of our lives. These rules most often take into account a very limited number of environmental factors, sometimes as few as one.

In addition to intuition, we have the uncanny ability to size up a situation by matching patterns in our current environment, to those we have seen in the past. Early researchers in the field of artificial intelligence, made certain to imbue their programming languages with similar capabilities. This mental facility allows us to identify things that are out of place (anomalies), or events which should be expected, but have not yet occurred. Pattern matching also allows us to perceive and recognize cues that help us to properly diagnose a situation, so we can know if we are on the right track. We also have the ability to infer new knowledge, from unrelated facts, even if some information is missing or incomplete. Within our minds we can simulate a sequence of events, starting at an existing state and leading to a specific goal or outcome. We have unlimited storage capacity inside of our brain, as well as the ability to compress many years of learning and domain specific knowledge, into a body of facts and rules that can be searched in a split second. All told one would believe that the human mind is perfectly suited for making decisions.

Limitations to Human Cognitive Abilities

To describe only the positive aspects of our decision making capabilities, without telling of our shortcomings would paint a very unfair picture. Although our minds have been gifted with the abilities told above, we have some fatal flaws as well, that hinder our ability to make good decisions. These flaws are inextricable from our mental processes, and are unfortunately part of the essence of human nature itself. First of all, we oftentimes incorrectly interpret significant events. For example if a manager in a large corporation is terminated, some may assume that it is for missing deadlines or going

over budget, when in reality very few employees are actually dismissed for these reasons. A second series of flaws arise from over confidence in our expertise and hence the decisions we make. As human beings we even go so far as to find supporting evidence for our decisions, even to the point of ignoring critical indicators. For example, during the Iran – Iraq war of the nineteen eighties the Vincennes, a United States navy Aegis class cruiser stationed in the Persian Gulf, mistakenly shot down a passenger airliner. When the incident was analyzed in retrospect, it became evident that two members of the command and control team, incorrectly stated to the commander of the ship that an unidentified object on the radar screen was descending, when in fact it was actually climbing, even though on board instrumentation indicated otherwise. It was deemed at the board of inquiry that the men based their statements on their own pre-disposed beliefs that the object in motion was in fact hostile, and the ship was about to be attacked, completely ignoring legitimate environmental cues. People are emotional. This tends to make us biased in our thinking and very judgmental. Our innate ability to match patterns, which benefits us greatly in assessing situations, as well as our tendency for our perceptions to become distorted by our own expectations, combine to form negative stereotypes within our minds. These stereotypes can easily lead us to make poor decisions.

Other well known human emotions such as greed and fear, can also dramatically affect our ability to make sound decisions, the sub-prime financial crises is just one timely example of this. Finally, human beings are not very good at determining the importance of, and correlating, a large number of diverse data elements. A long time ago Bell Labs researchers concluded that people can only effectively work with about seven different pieces of data at one time, which is why phone numbers in the United States have those many digits. As a result human beings are not very good at understanding the underlying relationships between large amounts of data. In today's complex world though, decisions oftentimes need to be based on over one hundred different pieces of seemingly unrelated information.

Automated Decision Making Systems

Making a decision involves selecting a course of action to follow. In a corporate environment the decision may be geared towards profitability, or a long term strategy of the organization. On the battlefield it may mean deciding where, when and how to attack an enemy. It involves the analysis of a situation, including the actions of other entities that may be involved, be they customers, competitors or enemy troops. During the process, relevant information needs to be extracted from a database or warehouse, or situation awareness, then translated and cleansed. Human knowledge needs to be discovered and embodied within an appropriate repository. Mathematical formulas need to be applied to the data. After alternative choices have been evaluated, a course of action is selected and initiated.

Driving Forces

Business Survey Results

A 2004 Opinion Research Corporation survey found that ninety percent of executives believed that front line operational decisions affected their company's bottom line. More than half of the companies interviewed though, had automated twenty-five percent or less of these decisions. In the same survey about fifty percent of the respondents admitted that it took months to change the decision logic that was embedded within their existing main line of business applications. In a 2004 Teredata survey, seventy percent of the respondents stated that poor decision making was a serious problem for their business. In an October 2006 PricewaterhouseCoopers Management Barometer survey, eighty-four percent of executives cited their inability to mine and interpret data as the single highest ranking obstacle to achieving real business value.

Core Applications No Longer Provide a Competitive Advantage

The first sixty years of the computer revolution have mainly focused on the automation of rote manual tasks. This includes such business processes as facilitating and managing transactions, enterprise resource planning, finance and accounting, and customer relationship management. While providing software consulting services over the past twelve years it has become evident that client management, policy administration, billing, and claims processing systems developed at different insurance carriers are amazingly similar, in terms of features and functionality. But what really gives one firm a competitive edge over the other, is their ability to accurately assess risk exposure, set rates to maximize profitability on a per client basis, provide customers with the product choices that they want, and identify fraudulent claims. In essence, what we would call the traditional "core line" of business applications, no longer play a part in giving a business a competitive edge. What really gives one company an advantage over another, is its proprietary knowledge of the industry, the marketplace in which it operates, as well as a detailed understanding of their customer's behavior.

Gap in Traditional Business Intelligence Applications

The term "business intelligence" was first coined in an article appearing in the IBM Journal back in 1958. The author was H.P. Luhn. In his writings he chose to refer to the Webster's collegiate dictionary definition of intelligence: "the ability to apprehend the interrelationships of presented facts in such a way as to guide action towards a desired goal." In 1989 a Gartner analyst appropriated the phrase and applied it to a classification of software then known as decision support systems (DSS). These applications, originating in the late nineteen eighties, have evolved over the last twenty years into the widely recognized business intelligence suites of today. These applications provide corporations with sophisticated reporting capabilities, dashboards by which a manager can easily visualize all of the key performance indicators within their department or division, as well as the means to dice and slice information in a myriad of different ways.

For the most part though, these systems still fulfill their original role, that of decision support, and fail to live up to their moniker, which infers there is some type of “intelligence”.

Underlying Technologies

The two core technologies which underlie current implementations of automated decision making systems are predictive analytics and a business rules engine. Predictive analytics encompasses a variety of techniques taken from statistics and data mining that analyze current and historical data in order to foretell future events. Data mining focuses on uncovering hidden patterns, and relationships, that are contained within information. The endeavor is sometimes referred to as “knowledge discovery in data” or KDD for short. The process of data mining entails the construction of a “model”, which may include a set of rules, or algorithms, which when applied to input data produces a “score” or result. The primary types of models are classification, estimation, and prediction. Classification looks at the attributes of an object, and then assigns it to a specific category, based upon the analysis of an existing data set in which the values of that particular object’s class are already known. While classification involves discrete values, such as yes or no, or the type of product, estimation relates to continuous values, such as a credit card balance or yearly gain. Prediction is a variation of classification or estimation, except that the results lie in the future.

In 2000 the Business Rules Group, suggested the following definition “A business rule is a statement that defines or constrains some aspect of a business. It is intended to assert business structure or to control or influence the behavior of the business”. A business rules engine consists of software infrastructure that allows the execution of business rules. It usually includes a language by which to write rules, as well as a repository in which to store and manage them. Some rule engines only allow for the sequential execution of rules, while others support the ability to make logical inferences. A number of vendors who manufacture business rules engines have incorporated the industry standard Reté algorithm, which was invented by Charles Forgy in the late 1970’s while working on his Ph.D. at Carnegie Mellon. Current automated decision making systems are based upon one or both of the aforementioned technologies.

Benefits of Decision Automation

The benefits of automating human decisions are manifold. The first of which is an increase in the speed at which a decision can be made. The technology can reduce the time it takes to underwrite and bind an insurance policy, or approve a lease application, from a few days down to a number of seconds. Second of all, the accuracy of the decision itself can be significantly improved. For example, a decision service that is part of a system to diagnose and treat medical conditions, can incorporate rules about drug interactions and adverse clinical outcomes, along with predictive models linking unexpected results to genetic variants contained within a particular patient’s DNA. This type of system could eliminate thousands of unnecessary deaths every year. Third, this

technology will allow machines to make unbiased decisions which are consistent throughout all lines of business and customer touch points within a corporation, or within every unit in a military fighting force. As stated before, one fault of the human thought process is that we are judgmental by nature, and as a result, rely upon stereotypical views of our environment when making a decision. Automated decisions will completely eliminate human bias. Just as important, automating human decisions will effectively lower the human capital costs associated with making a decision. Lastly the decisions will always be compliant, both with organizational strategies as well as government or military regulations.

Why Expert Systems Failed

One can never predict the future without first understanding what happened in the past. During the nineteen eighties there was a gold rush to commercialize expert system technology. Startup companies in the field had gotten hundreds of millions of dollars worth of venture capital funding, but what was predicted to be an industry with over four billion dollars a year in sales never materialized. By 1990, yearly industry revenue estimates of six hundred million dollars were deemed generous. Although some highly visible projects such as the Authorizers Assistant, developed by American Express, and the expert VAX configuration tool, called XCON, were publicly discussed as symbols of the technologies prowess, the overall consensus was that AI did not live up to its expectations.

In 1995, T Grandon Gill, a professor at Florida Atlantic University in Boca Raton, wrote a paper entitled "Early Expert Systems: Where Are They Now?". During the course of research he studied about seventy or so expert system projects that had started during the 1980's, covering both shrink wrapped software packages, as well as in house applications developed by American companies. His findings indicated that between 1987 and 1992 most of the systems studied fell into disuse, or were completely abandoned. What was interesting to note was that in the majority of cases, the relatively short software life cycle was not attributed to any technical deficiency or economic failure of the system, but rather due to a lack of acceptance by users, loss of key personnel, or shifts in organizational priorities.

But the most insightful finding within the paper plainly identifies the single most disconcerting obstacle for the acceptance of any expert system, human nature itself. As Mr. Gill aptly stated "Conventional systems such as accounting applications, often automate such tasks that are already routine and performed at a low level in the organization. Expert systems in contrast, tend to be applied to the types of tasks performed by individuals with greater skills and higher positions in the organization. As a consequence, potential users may be particularly sensitive to a technology they perceive as intruding on their task domain, and from a purely pragmatic standpoint, may be in a strong position to resist such technology."

First Generation Systems

As with any new technology there is an evolution from its initial introduction until its retirement from service. Think for a moment of the invention of the automobile around eighteen ninety, and the course grained, faded black and white photos of what looked like a horse drawn carriage of the era that had been fitted with an early version of a gasoline powered, internal combustion engine. Contrast those images to today's sleek, aerodynamic sports cars with leather bucket seats, climate control, and three hundred horsepower engines. In time we will experience a similar product life cycle with automated decision making systems.

In terms of their technical maturity, current automated decision making systems would be loosely classified as first generation. For the most part they automate high volume operational decisions. From a tool, development, and implementation perspective, lessons were well learned from the failure of earlier expert systems. The rules engines of today are much easier to use than the complex AI languages of the past. Most importantly they do not require a painstaking process by which to extract business knowledge from subject matter experts. Users who had previously resented the codification of their domain knowledge now routinely participate in writing and maintaining business rules through web browser based front ends. Data mining toolkits no longer require a PhD in statistics or mathematics to use effectively, but rather are targeted more towards marketing professionals. KDD tools now have graphical user interfaces, including wizards, visual editors to diagram workflows, and easy to use components that can rank the relative importance of incoming data elements and compare the reliability of different models after they have been built. What is also important to note, which has helped with their overall acceptance, is the fact that automated decision making systems are no longer deployed in a stand alone fashion, but rather have been fully integrated into core business processes and applications. In addition, they not only make a decision, but have been imbued with the ability to immediately carry it out.

Current Implementations

Currently automated decision making systems are being used across industry boundaries and departmental lines. In the hotel industry systems of this nature handle price optimization for room rates. One casino operator, Harrah's, pioneered the use of a customer loyalty program in conjunction with price optimization in order to offer room rates to individual customers based upon their anticipated long term value to the business. This form of "one on one" marketing is an excellent use of what the technology has to offer at this point. In the insurance industry policy underwriting decisions in many firms have been, or are in the process, of being automated. In the financial services industry trading and underwriting processes have been automated. In any business environment, if you can reduce the amount of time it takes to complete a core process, down from several days to a number of seconds, there is an incredible impact on the on a companies bottom line.

In the internet age, when a potential customer can visit a number of different websites within a few minutes, the ability to close the “sales loop” in a matter of seconds is crucial from the business perspective. In the late nineteen nineties I consulted on such a project at Compaq Financial, the leasing arm of the computer maker. The goal was to build a system which would automate the credit evaluation and lease approval process, bringing it down from several days to just under two minutes. The system communicated directly with Dunn & Bradstreet computers in order to obtain credit data on a small business. Given that information, a credit score would be calculated, and a number of rulesets would fire in order to determine that the application was not fraudulent and that the lease should be approved, sometimes with additional requirements. The system supported a completely automatic approval process, with the system itself defining contingencies such as a personal guarantee by the lessee, or a larger security deposit than originally estimated. The system also had facilities for manual intervention in the outlying cases in which it was warranted. The banking industry, over the last fifteen or so years, has developed similar systems for the real-time approval of mortgages, home equity loans, and personal charge card purchases.

Some of the most publicly visible uses of automated decision making systems are internet based product recommendation engines. One example would be Netflix's electronic facility to suggest video entertainment based upon your previous account activity, user profile, and reviews of the movies you have already rented. Another would be Amazon.com's facility to recommend titles based upon the books you have previously bought. Although most uses of this type of technology are currently targeted towards the retail industry, a product recommendation engine could easily be used by a financial services firm, insurance company, pharmaceutical manufacturer, or travel related company.

Limitations of Existing Systems

Although first generation of automated decision making systems have been successful overall, given the business and technical goals that they have been designed to achieve, they are less than perfect. In addition, most of the vendors providing rule engine and data mining products to this emerging market, are optimistically cautious in their claims of decision automation. The applications which are readily discussed within marketing literature and trade publications tend to focus on high volume operational decisions, which need to be made in a relatively short period of time. The importance of these decisions as a whole is great, but the value (or loss) of any given, potentially incorrect decision, is relatively small. Case studies or published writings about higher impact decisions, such as recommending a drug therapy, given a clinical diagnosis and patient history, are not easily found within publicly available writings on rules engines or data mining tools.

To gain a better understanding on the potential downside of first generation of decision automation technology, we can look at the financial services industry and the subprime mortgage crises. The number of houses currently in foreclosure in the United States is at an all time high. A percentage of these mortgages were approved, using some type of automated decision making system. Plain common sense would lead us to conclude that many of the borrowers who were approved for these loans, taking into account that they

had “spotty” credit histories at best, would never be able to make the monthly payments over the life of the mortgage.

After the loans closed they were sold off to investment banks who underwrote the securitization of these assets. The financial instruments that were created were called collateralized debt obligations, or CDO’s for short. Many of the investment banks chose to keep a portion of these financial instruments, called tranches, for themselves, selling off what they considered to be more risky slices of the pie. This is where financial modeling techniques began to fall short, as the market for CDO’s was illiquid and potentially subject to huge price discounting. In addition, many of the pricing and risk models did not take into account housing values themselves, the stock price of the mortgage company, the fact that the borrower produced no documentation on their income, was given a one percent “teaser” rate, or was guaranteed money back at closing.

In October of 2006 the ABX, which is a measure of the risk of owning bonds backed by subprime mortgages, rose thirty percent over its position two months prior. This was a sign of trouble which well designed financial models should have detected. During this time some savvy traders at a number of hedge funds, as well as investment banks, knew that they should invest in credit default swaps, in order to offset any potential losses should the value of their CDO positions fall. This trading strategy brought to light a noted shortcoming of many automated trading systems, the basic premise of “hedging one’s bets”. Unfortunately this key concept was not enforced by the automated trading systems at most firms holding substantial CDO positions. As it turned out, the risk and pricing models that were built failed to identify the potential that these CDO’s would fall dramatically in value.

In an article entitled “Can Financial Models Prevent CDO Problems”, published in September of 2007 by Wall Street & Technology, Lyle Minton of Point Clear Capital Management aptly stated “You can get lost in your financial modeling...So many people model up and have well thought out trading positions based upon historic valuations without thinking about the common sense aspect of when you get in true market stress, liquidity dries up so you can’t get the execution you thought and historic correlations fall apart”.

Other shortcomings of first generation systems lie not in the actual approach to using the business rule engine, quantitative analysis tool, or algorithm itself, but rather with the underlying infrastructure upon which these solutions are constructed, and eventually deployed. Existing enterprise software infrastructures, upon which today’s automated decision making systems are built, were originally designed to support transaction based systems. This type of infrastructure is primarily dependent upon means for data input and output, such as databases, message buses, printers, display screens, or network communications. The programming model traditionally used with this type of infrastructure is sequential, meaning an application would have a beginning, middle, and an end, assuming of course there would be loops and decision points within a flowchart. Executing the program would lead to an eventual success or failure, amounting to a commit or rollback of information within the database. From the system perspective this transaction might correspond to the sale of goods.

Automated decision making systems on the other hand are completely different. They are comprised of complex mathematical formulas that build decision trees, train neural networks, and execute logistical regressions. Then there are algorithms like Reté, which is responsible for the pattern matching which takes place inside of the business rules engine. The process of pattern matching allows a piece of software to simulate human cognitive capabilities, such as is required when making logical inferences. The applications usually are not executed sequentially, as rules are declarative (non-procedural) by their very nature. The perfect environment in which to develop and deploy this type of software would be one that supports high performance parallel task execution, application virtualization, a GRID comprised of commodity servers, and a service oriented architecture.

Conclusion

We live in a world filled with uncertainty in which newer business models are replacing those which proved successful throughout the industrial age. In a similar light age old theories of rational thought, relating to how human beings make decisions, have been disproved as well. Psychologists now realize that intuition and common sense reasoning play a critical role in all of the decisions we make. But even with our highly evolved cognitive capabilities, humans have some significant flaws which prevent us from consistently making fast, accurate and unbiased decisions.

After over sixty years of software development, almost all rote manual tasks within an organization have been automated. Further efforts in these areas will result in little or no financial gain. On the other hand, the only true competitive edge that any organization has is the knowledge it possesses about its customers, the marketplace in which it operates, and the products and services it provides.

Traditional business intelligence applications, although providing valuable insights into the operations of an enterprise via reporting, dashboards, and other similar tools, fall short of making any firm recommendations. Automated decision making systems utilize business rules engines and predictive analytics in order to choose and execute the best course of action.

The benefits of using automated decision making systems include faster, more consistent, and more accurate decisions, which are free of human bias, follow organizational strategies, and abide by government regulations. First generation systems are the direct descendents of expert systems that were developed within the 1980's. But they have been built using more user friendly tools, with the full support of the organizations in which they are being deployed.

For the most part first generation systems to automate human decisions have met their design goals, and have been well received from a business perspective. But there is room for improvement, as the recent subprime financial crises has revealed. Using current approaches it is difficult to build a financial model which takes into account such

factors as fear or greed. And although a rules engine can embody some knowledge about customer relationships, a specific marketplace, industry, and product offerings, current system implementations lack any facility to simulate human cognitive abilities that are similar to common sense or intuition.

In 1997 a computer system named Deep Blue, defeated Gary Kasparov, a chess grand master. The computer was capable of two hundred million computations per second, far greater than any human. It could easily project out the next eleven moves, hence its name, while Gary Kasparov could only envision four or five. It was not any smarter than the grand master, although it undoubtedly contained a number of heuristics for playing chess. It simply out computed him. But if Deep Blue was rolled outside of the data center in which it is housed, the wind blew and clouds darkened, and then it began to rain, the machine would not have the common sense to go back inside. Within a matter of minutes Deep Blue would undoubtedly short circuit and shutdown. Ironically, coming in out of the rain is something everyone learns in childhood.

About Mobile Agent Technologies

Mobile Agent Technologies builds machines that think, those which exhibit human like cognitive behavior and can make core business decisions. In doing so we have created an integrated environment upon which to build intelligent applications called Einstein Enterprise. The platform integrates an inference engine, complex events processor, data mining toolkit, and knowledge repository, with our own patented mobile agent based distributed computing environment. We also incorporate our own proprietary methodology by which to imbue software systems with intuition and gut feelings, which we call Common Sense Reasoning™. To learn more about our products and services please contact us at 201-923-1006, or visit our website www.agentos.net.