Transformation: The Human Obstacles

Jens Pohl
Collaborative Agent Design Research Center
Cal Poly State University (Cal Poly)
San Luis Obispo, California, USA

Abstract

This paper addresses the human aspects of change at a time when the rapid wide-spread implementation of a new infostructure has become the cornerstone of national security policies. The author explores the behavioral characteristics of human beings from a point of view that is not based on inter-human comparisons, but on an examination of the innate physical, biological, intellectual, emotional, and social characteristics of our species. It is argued that the experience-based nature of the human cognitive system greatly reduces our ability to adapt to changes in our environment, exploit opportunities, and create new knowledge. A fundamental biological survival instinct drives the human being to seek a level of certainty and security that is irreconcilable with a continuously changing and largely unpredictable environment. The symptoms of this distinctly human struggle manifest themselves in a strong resistance to change and an emotional aversion to experimentation and risk taking.

As an example that is currently of critical importance, the paper discusses the reluctance of human beings to accept computers as any more than unintelligent computational machines. Focusing on the collection and analysis of intelligence data in a post-Sep.11 (2001) world that is subjected to the constant threat of terrorism, the author points out the human obstacles that have delayed the necessary paradigm shift from data-centric to information-centric software. By presenting an assessment of human weaknesses and strengths the paper attempts to provide some useful insights into human behavioral characteristics that will tend to create a great deal of tension during this critical period of accelerated change.

It is the humble hope of the author that the notions expressed in this paper will provide cause for further thought and deliberation, and may contribute in some useful manner to diffusing the intense human intellectual and emotional struggles that characterize periods of rapid change. Specifically, the author is greatly concerned that these human struggles may further delay the utilization of higher level computer capabilities, at a time when such capabilities are so desperately needed.

Compelling reasons for change

There have been periods of rapid and profound change throughout the history of the world. More often than not, and certainly in recent times, the precipitating factors have been technological and/or political in nature. Sometimes these factors have gained momentum over time in a cumulative manner such as the French Revolution in the 18th Century, and at other times they have descended on society more abruptly. The terrorist attacks on the USA that occurred on September 11, 2001 are an example of the latter. In either case such periods of change have typically been accompanied by a great deal of human tension.

It is the purpose of this paper to explore some of the underlying reasons for the tensions that appear to inevitably accompany periods of rapid change in society. Focusing on the possible repercussions of the terrorist events of September 11 (2001) this exploration is timely for several reasons. Firstly, the threats posed by potential follow-on terrorist activities are both very serious and immediate. In fact, they are considered so dangerous that the US Government has found it necessary to initiate a degree of mobilization and reorganization that is unprecedented since World War II. Secondly, the current vulnerability of the US to particular kinds of terrorist attacks calls for urgent remedies. The development and implementation of such remedies must not be unduly obfuscated and delayed by lack of cooperation due to human resistance to change. While it may not be possible to altogether eliminate such inherently human characteristics, a better understanding of the nature and symptoms of the human disposition may allow individuals and groups to contribute to solutions in a more effective manner. Thirdly, during a time of rapid change there is a need for leadership with vision and perspective. Both of these highly desirable attributes are based on insight into human nature.

Finally, at least the short term remedies to the prevalent terrorist threat are technology dependent. They require us human beings to reevaluate the very essence of our relationship with computer-based systems. To date we have shown a great deal of resistance to the notion that a computer is any more than a dumb electronic machine. The concepts of computer intelligence and human-computer collaboration are not generally compatible with our human view of the world. Yet, such computer-based capabilities have been demonstrated as being feasible, have been available for exploitation for more than a decade, and probably offer the most useful and immediate solution approach to combating terrorism.

Humans are *situated* in their environment

It would not be unreasonable to postulate that we human beings tend to have a rather inflated view of ourselves, in terms of our capabilities and our relationship with anything that is not human. The fact is that we have little to compare ourselves to apart from other human beings and lower forms of life. Whenever we assess our capabilities we do so in comparison with our colleagues. Arguably, this gives us a very biased view of ourselves and makes us vulnerable to a self-centered interpretation of our environment in general and more specifically the events that occur within that environment. In short, we have a tendency to believe that everything revolves around us.

Clearly, we are ***situated*** in our environment not only in terms of our physical existence but also in terms of our psychological needs and understanding of ourselves (Brooks 1990). We depend on our surroundings for both our mental and physical well being and stability. Consequently, we view with a great deal of anxiety and discomfort anything that threatens to separate us from our environment, or comes between us and our familiar surroundings.

This extreme form of ***situatedness*** is a direct outcome of the evolutionary core of our existence. The notion of evolution presupposes an incremental development process within an environment that represents both the stimulation for evolution and the context within which that evolution takes place. It follows, firstly, that the stimulation must always precede the incremental evolution that invariably follows. In this respect we human beings are naturally reactive, rather than proactive. Secondly, while we voluntarily and involuntarily continuously adapt to our environment, through this evolutionary adaptation process we also influence and therefore change our environment. Thirdly, our evolution is a rather slow process. We would certainly expect this to be the case in a biological sense. The agents of evolution such as mutation, imitation, exploration, and credit assignment, must work through countless steps of trial and error and depend on a multitude of events to achieve even the smallest biological change (Waldrop 1992, Kauffman 1992, Holland 1995, Pohl 1999).

In comparison to biological evolution our brain and cognitive system appears to be capable of adapting to change at a somewhat faster rate. Whereas biological evolution proceeds over time periods measured in millenniums, the evolution of our perception and understanding of the environment in which we exist tends to extend over generational time periods. However, while our cognitive evolution is of orders faster than our biological evolution it is still quite slow in comparison with the actual rate of change that can occur in our environment.

Human resistance to change

Clearly, at least in the short term, the experience-based nature of our cognitive system creates a general resistance to change. The latter is exacerbated by a very strong survival instinct. Driven by the desire to survive at all costs we hang onto our past experience as an insurance. In this respect much of the confidence that we have in being able to meet the challenges of the future rest on our performance in having met the challenges of the past (i.e., our success in solving past problems). We therefore tend to cling to the false belief that the methods we have used successfully in the past will be successful in the future, even though the conditions may have changed. As a corollary, from an emotional viewpoint we are inclined to perceive (at least subconsciously) any venture into new and unknown territory as a potential devaluation of our existing (i.e., past) experience.

This absolute faith in and adherence to our experience manifests itself in several human behavioral characteristics that could be termed limitations. First among these limitations is a strong aversion to change. Typically, we change only subject to evidence that failure to change will threaten our current existence in a significant way. An example is the rather slow transition from data-centric to information-centric computer software. Although the digital computer was originally conceived as a very fast computational machine capable of reducing the time required for the solution of large numbers of mathematical equations from days to seconds, it soon emerged as a data storage and processing facility. This was mainly due to the need for record keeping accelerated by the growth of commerce and industry driven by major improvements in the ability to travel and communicate over long distances. As a result new opportunities for interaction, leading to cooperation, and eventually collaboration, presented themselves. As the intensity of these activities and the tempo of daily life increased so also did the competition among the human players. However, it did not occur to these players for at least two decades that the functions of the computer could extend beyond the rote storage and processing of data to the representation of information as a basis for automatic reasoning capabilities.

Prior to the events of September 11 (2001) the gradual realization that human-computer interaction could be raised to the level of meaningful collaboration came not as a result of creative discovery, but because the requirement of interpreting the vast amount of computer-stored data simply outstripped the availability of human resources. In other words, it was not the opportunity for using computers in this far more useful role, but the necessity of dealing with an overwhelming volume of data that was gradually persuading computer users to elevate data-processing to information representation in support of automatic reasoning capabilities. Subsequent to September 11 (2001), the absolute necessity of automating at least the lower levels of intelligence gathering and analysis has begun to accelerate the transition from persuasion to conviction. Driven by the realization that the US can no longer afford to depend on the mostly manual processing of intelligence data, key government officials responsible for implementing a vastly improved infostructure have begun to seriously pursue an information-centric software architecture (Cooper 2002).

A second limitation is our apparent inability to resist the temptation of applying old and tried methods to new situations, even though the characteristics of the new situation are actually quite unlike the situations in which the existing methods were found to be useful. This typically casts us into an involuntary experimental role, in which we learn from our initial failures. Examples abound, ranging from the development of new materials (e.g., the flawed introduction of plastics as a structural building material in the 1950s) to the reluctance of the military to change their intelligence gathering and war fighting strategies long after the conclusion of the Cold War era in the 1990s (Wood 2001).

A third limitation is our tendency to view new incremental solutions as final comprehensive solutions. A well known example of such a problem situation was the insistence of astronomers from the 2nd to the 15th Century, despite mounting evidence to the contrary, that the heavenly bodies revolve in perfect circular paths around the Earth (Taylor 1949, 108-129). This forced astronomers to progressively modify an increasingly complex geometric model of concentric circles revolving at different speeds and on different axes to reproduce the apparently erratic movement of the planets when viewed from Earth. Neither the current scientific paradigm nor the religious dogma of the church interwoven within the social environment allowed the increasingly flawed conceptual solution of Ptolemaic epicycles to be discarded. Despite the obviously extreme nature of this historical example, it is worthy of mention because it clearly demonstrates how vulnerable the rational side of the human cognitive system is to emotional influences (Pohl et al.1997, 10-11).

The digital transformation

Over the past 30 years the principal advances in digital technology have been related to the miniaturization of electronic components, increases in the power of computer hardware, and the connectivity provided by communication systems. While these advances have been dramatic in terms of continuous increases in computational speed, memory, storage capacity, and decreases in cost, the intellectual utilization and exploitation of these capabilities has been less than startling. The latter is the province of the software that executes on the hardware and allows human users to accomplish tasks.

It can be argued that our human view of computer software has been shortsighted in respect to two popular notions: firstly, that data and information are essentially synonymous terms; and, secondly, that computer intelligence is largely a misnomer because computers are machines. Neither of these notions is accurate. While we human beings are able to convert data (i.e., numbers and words without relationships) automatically into information due to the experience (i.e., context) that is held in our cognitive system, computers do not have the equivalent of a human cognitive system and therefore store data simply as the numbers and words that are entered into the computer. For a computer to interpret data it requires an information structure that provides at least some level of context. This can be accomplished utilizing an ontology of objects with characteristics and relationships (Pohl 2001, Uschold and Gruninger 1996, Gomez-Perez 1994). In this way it is possible to create a virtual version of real world situations in computer software.

For example, let us assume that we wish to represent a component of a building such as a conference room in the computer. Until recently, in a ***data-centric*** software environment, we would have treated the conference room as a three-dimensional geometric entity that can be described in terms of points (i.e., x-y-z coordinates), lines, or surfaces. While this may be satisfactory for displaying different internal views of the building space and even generating animated 'walk-through' sequences, it does not provide a basis for the computer to reason about any aspect of the space, such as that a conference room must have a door for it to be usable. To provide the computer with such a reasoning capability the particular entity, in this case the conference room, must be represented in the computer as an information structure that constitutes the context of a building. This can be achieved quite easily by storing in the computer the word building and associating this word with some characteristics such as: a building is a physical object; it is made of material; it has height, width and length; consists of one or more floors; has spaces on floors; and so on. Then further defining spaces with characteristics such as: enclosed by walls, floor and ceiling; with walls having at least one opening referred to as a door; and so on.

In such an ***information-centric*** software environment the same conference room would be stored in the computer as part of the 'building' information structure (i.e., ontology), allowing the following more intelligent user-computer collaboration:

Computer user: I would like to represent a component of a 'building'.

Computer software: Loads its stored 'building' ontology into memory.
Asks user "What kind of a 'building' component?"

Computer user: A 'space' of type 'conference room'.

Computer software: For how many persons?

Computer user: Up to 16 persons.

Computer software: Suggested space size is: 16 ft (length), 14 ft (width), 8 ft (height).
Suggested furniture: 6 ft by 3 ft table, 16 chairs, screen, white board.
Other features: There must be at least one door.

As can be seen from this user-computer interaction, the computer software by virtue of its internal information structure has some understanding of the meaning of a building within the context of its characteristics and the relationships of its components (i.e., floors, spaces, walls, openings, and furniture). This endows the computer software with the ability to collaborate and assist the user by reasoning about the relationships between the data entered by the user and the context contained in the simple information representation provided by the 'building' ontology.

The potential impact of this kind of computer-based reasoning capability on current data-based intelligence collection and analysis practices is profound, to say the least. For example, let us assume that during a typical four-week period the following four messages (Fig.1) have appeared among the hundreds of thousands of data items that were collected by the various intelligence agencies: (message A) terrorist warning from Egyptian counterintelligence agency; (message B) suspected Middle-East terrorist apprehended on entry into US; (message C) meeting between suspected terrorist operatives at O'Hare Airport in Chicago; and, (message D) explosives theft in Chicago. Although the messages come from different sources and are collected by different intelligence agencies, they are automatically time-stamped and fed into a message network (Fig.2) that is presumably accessible to all such agencies.

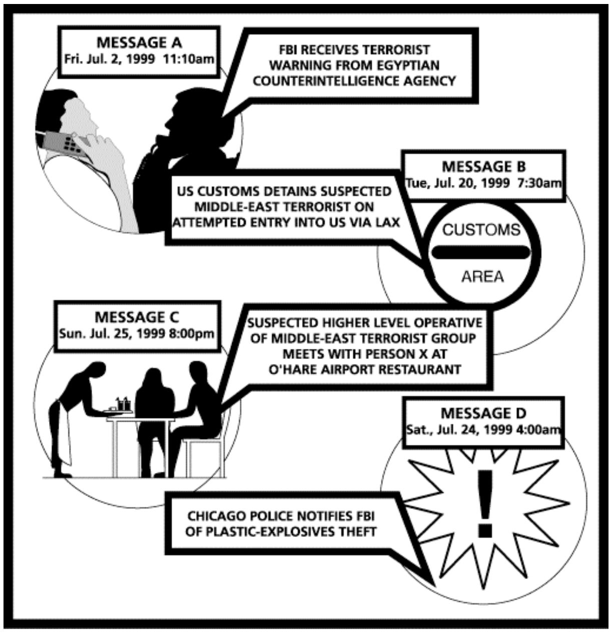

Fig.1: Four typical intelligence messages
(***data-centric*** software environment)

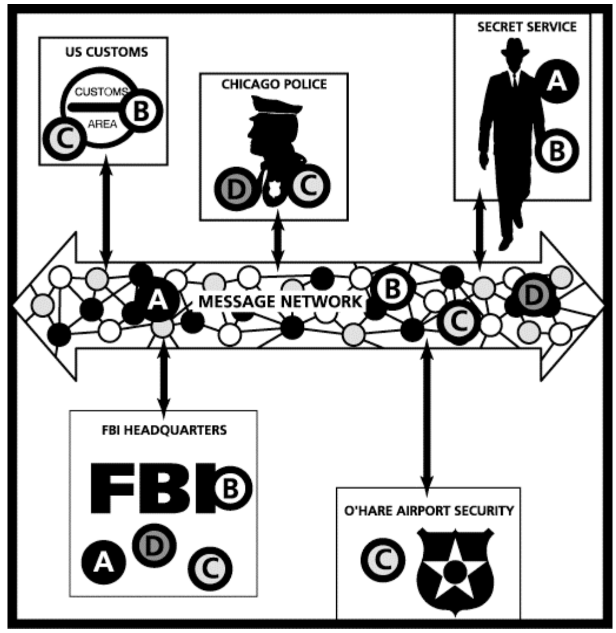

Fig.2: Manual processing of messages
(***data-centric*** software environment)

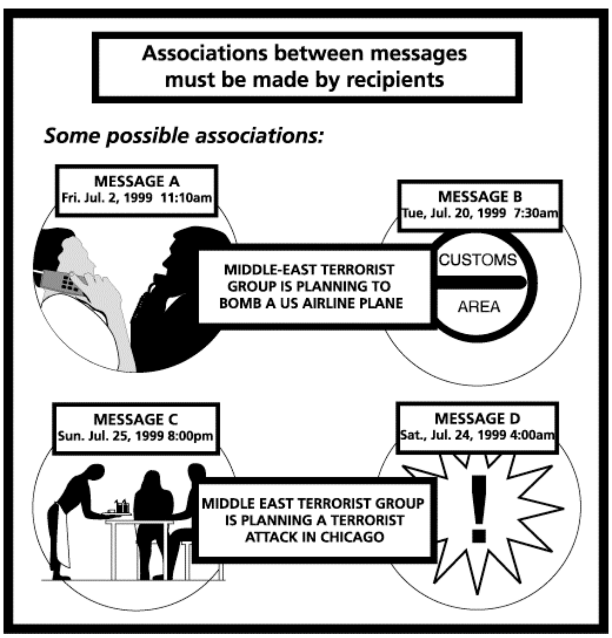

Fig.3: Manual processing of associations
(***data-centric*** software environment)

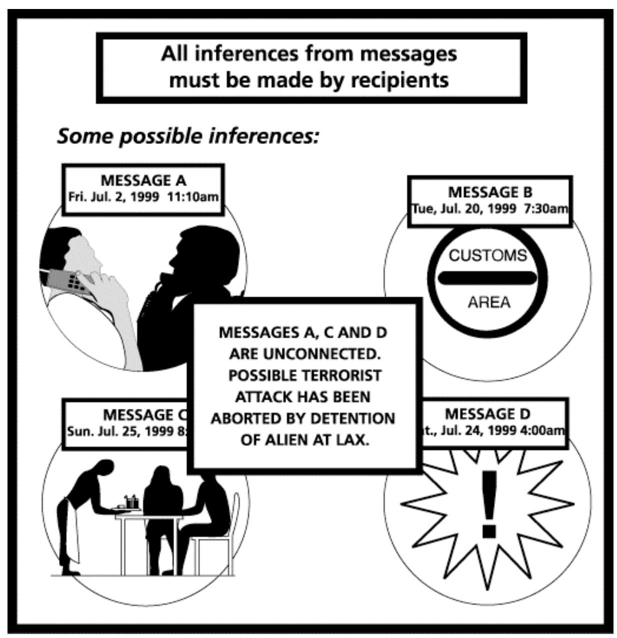

Fig.4: Manual processing of inferences
(***data-centric*** software environment)

There are several serious problems with this method of collecting and subsequently analyzing intelligence data in the current ***data-centric*** software environment. Firstly, virtually all of these raw data items have to be analyzed and evaluated by human operators. The absence of relationships (to provide context) prevents any really useful automatic filtering to be implemented. Only limited data-processing methods such as keyword searches and indexing, can be employed. This places an enormous burden on limited human resources. As a result valuable human intelligence personnel are literally burned out at the lowest level of intelligence analysis, leading to a tendency for the data collection activity to be restricted. Secondly, there is no guarantee that all agencies will receive all four of the data items. Since any associations between these and other messages have to be made by human analysts (Figs.3 and 4) there is a good chance that some of the possible associations will be either overlooked, or not pursued for lack of available manpower. Yet these associations are critical for the transformation of data into information and the detection of patterns through appropriate inferences. Even if the available human resources allow first level associations to be established, it is unlikely that many second and higher level associations will be developed due to human labor limitations and time constraints.

Nowhere are the shortcomings of a data-centric software environment, in which the computer-based systems have no understanding of what is being processed and virtually all interpretation tasks must be performed by human operators, more apparent than in the intelligence community. While the information-centric building blocks that would allow computer software to play a far more powerful role in intelligence analysis and evaluation processes have been available for at least two decades, the intelligence community like most other computer users has been reluctant to take advantage of these potential capabilities. Clearly, this reluctance is not based on technical obstacles but on the innately human resistance to change. Unfortunately, only a compelling reason such as the post-Sep.11 (2001) terrorist threat will provide the necessary incentive for the timely adoption of information-centric software design and implementation principles.

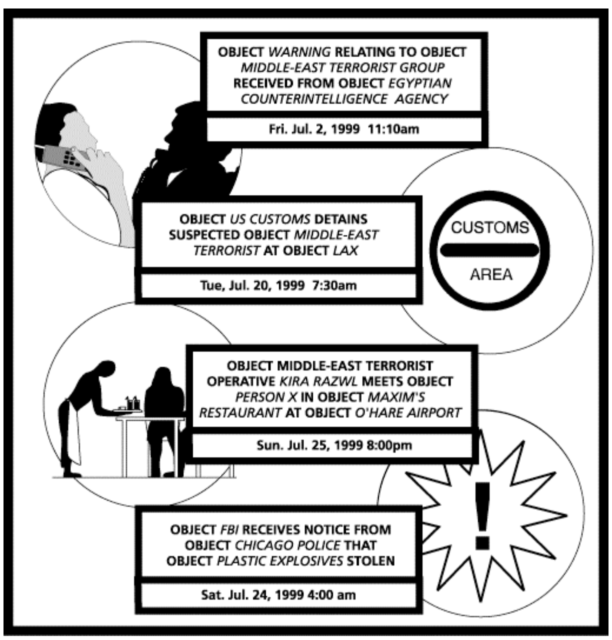

Fig.5: Four typical intelligence messages (***information-centric*** software environment)

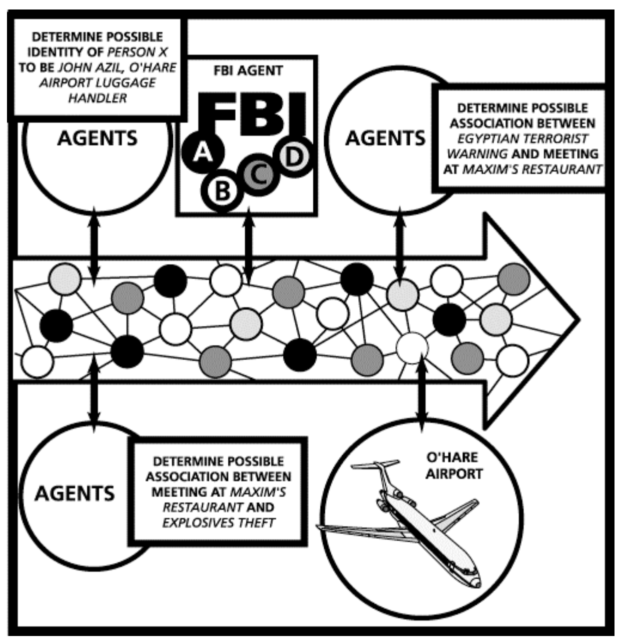

Fig.6: Automatic processing of messages (***information-centric*** software environment)

The benefits of such a paradigm shift will be immediate with startling results, even in the initial relatively primitive implementations. Returning to the previous example, the four typical intelligence messages would be processed very differently by ontology-based software that provides sufficient context for some degree of automatic reasoning by software agents. What is immediately obvious in Fig.5 is that key elements of the messages are treated as objects with attendant characteristics and relationships. For example in the first message, by being represented as objects within an information model the Egyptian Counterintelligence Agency, the Middle-East Terrorist Group, and the Warning itself, provide a rich context from which inferences can be drawn. This context includes not only the information that is already available about the Middle-East Terrorist Group and the Egyptian Counterintelligence Agency, but also the relationships among these entities and other entities (i.e., objects) represented in the information framework. This allows software agents (Fig.6) to automatically explore numerous associations and possible conclusions in parallel and without initial human assistance. As a direct consequence scarce human resources can be employed at higher levels of intelligence analysis and evaluation, after most of the tedious low level filtering and data correlation tasks have been accomplished.

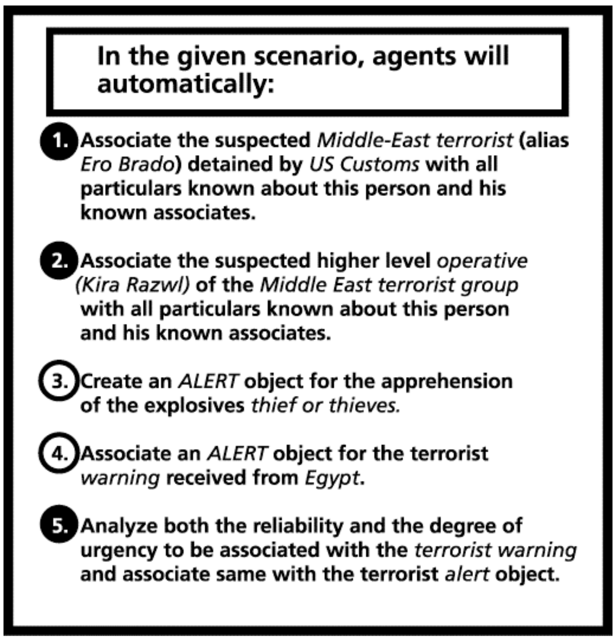

Fig.7: Automatic processing of associations (***information-centric*** software environment)

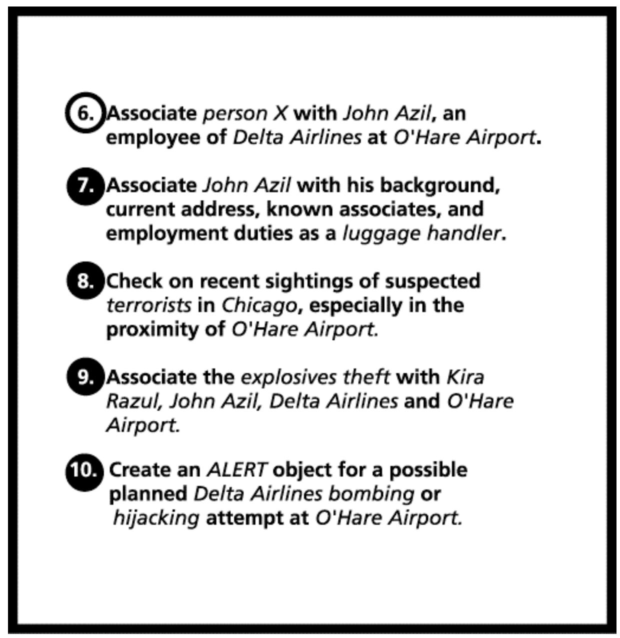

Fig.8: Automatic processing of inferences (***information-centric*** software environment)

A typical sample of the kinds of automatic analysis and reasoning that could be performed by software agents in the given example is shown in Figs.7 and 8. Not only would the agents be able to automatically access relevant databases to enrich the information environment, but they would also be able to correlate data sources, form associations, detect patterns, and explore possible conclusions. The number of software agents performing these tasks in parallel would vary depending on current needs and hardware capabilities. In other words, agents could be cloned by other agents, or even themselves, as demanded by the workload.

Following on from these examples of information-centric computer software, let us briefly reconsider the second statement posed at the beginning of this section: "…computer intelligence is largely a misnomer because computers are machines". Even if we are willing to accept the notion that computer software can have reasoning capabilities, does that equate to intelligence? From a human point of view, hardly. However, one could ask: Is human intelligence the only form of intelligence? Arguably there are different kinds of intelligence and also levels of intelligence. The lowest level of intelligence is remembering. Certainly, computers have the ability to remember, in fact, one could argue that they surpass humans quite easily in this area. Reasoning is a higher level of intelligence. Clearly, computer software can be designed with reasoning capabilities as long as there exists an underlying information structure that provides context. Learning is an even higher level of intelligence, yet elementary learning capabilities have been demonstrated in computer software for several decades (Forsyth 1989, Clancey et al. 1994). It would therefore appear that although computers are electronic machines they can function with some level of intelligence.

Symptoms and causes of human tension

Survival is certainly one of the basic biological instincts of all animals. The human manifestations of this instinct are readily apparent in a strong desire for absolute certainty and ***security*** in a changing and largely unpredictable environment. We have great difficulties coming to terms with the fact that many of the forces that influence our situation are beyond our individual influence and control. We endeavor to make long range predictions, even though time and again we see evidence that these predictions fall far off the mark. Not only do we insist on such predictions, but we also apply the most precise mathematical techniques in a misguided attempt to ensure accuracy. While developments in the field of mathematics were driven by the need for tools that could be applied to current endeavors, there has always been a correspondingly strong desire to apply these tools to the prediction of future outcomes and events. Elaborate mathematical models are often constructed to simulate the interaction of complex relationships over extended periods of time. When we are forced to explore why the real world behavior has diverged significantly from these simulations, we find that not only was the mathematical model simplistic and incorrect in the postulated behavior of individual variables but that many of the boundary assumptions were erroneous.

Other symptoms of the tension induced by the intrinsic human struggle with insecurity include: our insistence on applying only true and tried methods to current problem situations even though those situations may differ significantly from previous experiences (i.e., we simply find it very difficult to overcome our emotional aversion to experimentation and the risk of failure); our strong resistance to change (i.e., we typically have to be forced to change by an impending threat); our need to explain any currently unexplained phenomenon to eliminate uncertainty, even if we have to oversimplify the complex behavior of the phenomenon; and, our preference for ready-made solutions over tools (i.e., even though the tools would provide us with the flexibility of adjusting to real world conditions).

Another source of tension is the apparent desire to ***dominate***. It is a common trait in human beings to seek control and compliance, and to be intolerant to individual preferences. Typical examples range from: hierarchical organizational structures; unwillingness to share information; monopolies and unfair trade practices; orders without explanations; social classes and castes; dictatorships; regional invasions and wars; and, ethnic cleansing.

However during the past decade this source of tension has been mitigated by an increasing focus on the individual and the flattening of the traditionally favored hierarchical organizational structure into a web-like structure with a high degree of local autonomy. Driven mostly by the widespread availability of information technology, this has placed an unprecedented emphasis on knowledge and the value of human capital in a knowledge-based organization. As a result the notion of knowledge management has started to receive considerable attention in both government and corporate organizations. Simply stated knowledge management involves the effective acquisition, development and utilization of the human capital in an organization. The emphasis of this definition is on maximizing the contributions of the individual to the collective benefit of the organization. In this respect knowledge management serves primarily as a facilitating vehicle, with the objective of enabling the human and organizational capabilities for the benefit of the individual and the organization.

Through the distributed framework of leadership and communication made possible by the availability of information technology, knowledge management is able to execute its enabling role in several ways. Firstly, knowledge management recognizes that every member of the organization is a contributor and a potential decision maker. Therefore its methods are designed to emphasize the encouragement, cultivation (e.g., professional development), and motivation of the individual. Secondly, by definition, knowledge management emphasizes local autonomy and concurrent activities. This recognizes that leadership should be exercised through example, clarity and communication, and not through authority. Under these conditions the principal tools of leadership are the continuous analysis of feedback, the meticulous explanation and justification of intent and direction, and the maintenance of effective self-development opportunities throughout the organization. Thirdly, knowledge management tends to foster the formation of internal and external relationships, because the relationship capital of the organization becomes one of the most important catalysts for increasing the productivity of the organization.

Related to the desire for dominance is the inherent human behavioral characteristic of ***intolerance***. Generally speaking human beings seek conformity with currently accepted norms and tend to be intolerant of individual differences that deviate from these norms. Human society has always been pervaded with discrimination based on race, social position, and even personal appearance, as well as political and religious persuasion. In recent times democratic governments have found it necessary to pass anti-discrimination laws to, for example, safeguard the interests of ethnic minority groups, force consideration of the special needs of physically handicapped persons, and prevent the superficial profiling of criminal suspects by law enforcement authorities.

There is some doubt whether human beings are at all capable of altruistic behavior. The strong survival instinct leads to ***selfishness*** and behavior that is often detrimental to others. Certainly there is a pronounced human tendency to react aggressively to any perceived physical threat or psychological denigration.

Our deep-seated self-preservation instinct manifests itself in several additional ways. For example, we find it very difficult to view a situation from any perspective other than our own. This leads to many misunderstandings and unnecessary retaliatory actions at all levels of human interaction. It also constitutes a major obstacle to meaningful cooperation among individuals and successful negotiations among groups. On a more personal level, individuals find it difficult to admit mistakes for fear of losing status in their interactions with others, and perhaps even more perilously for fear of losing confidence. In this respect human beings find it very easy to invent reasons that allow them to blame others for their own misfortunes.

The inherent human resistance to change, discussed previously, is accompanied by a general ***complacency***. Relatively few human beings are consistently able to motivate themselves to actively identify and pursue new opportunities. The vast majority judiciously avoid voluntary intellectual explorations and appear to lack the confidence to take risks. As a rule, we prefer to blindly continue rather than question old habits and preferences. In summary, we are uncomfortable with change, generally frightened of the unknown, and likely to support the status quo unless extremely threatening circumstances exist.

Human limitations and weaknesses

Deeply embedded in the evolution of the human intellect is the rationalistic approach to problem solving. At face value this approach appears to be entirely sound. It suggests that problem solving should proceed in a logical sequence of clearly defined steps. One begins by defining the problem and then decomposing the defined problem into sub-problems. Decomposition appears to make a great deal of sense because the parts of a problem are intrinsically easier to solve than the whole problem. The reason for this is that the complexity of a problem is normally due to the nature and number of relationships among the elements of the problem and not due to the elements themselves. Decomposition allows us to temporarily neglect consideration of many of these relationships. However, this over-simplification of the problem is valid only as long as the problem remains in a decomposed state. As soon as we try to integrate the separate solutions of the parts into a solution of the whole the relationships that we so conveniently disregarded reappear and invalidate many if not most of our neatly packaged sub-solutions. We find to our consternation that the characteristics of a part of a problem situation considered in isolation are not necessarily similar (let alone the same) as the behavior of that part within the context of the whole problem.

Within the rationalistic paradigm this forces the human problem solver to repeat the decomposition process multiple times in a cyclic manner, shaping the sub-solutions to incrementally accommodate the influences of the various relationships until an apparently acceptable solution of the whole has been found. This process is laborious and in many complex and dynamic problem situations does not lead to a solution that should be considered as acceptable, but is nonetheless often adopted out of desperation.

Why have we human beings come to rely so heavily on this flawed approach to problem solving? The reasons are related primarily to the biological nature of our cognitive system, and secondarily to some of the tensions discussed in the previous section. While the biological basis of human cognition is massively parallel (i.e., millions of neurons and billions of connections, referred to as synapses) our conscious reasoning capabilities are largely sequential. The fact is that our short-term memory has a severely limited capacity of only a few chunks of data at any one time. Therefore, we can differentiate among only a small number of objects at any one point in time, even though we continuously move new data chunks from long-term memory into short-term memory. As a consequence we have great difficulty dealing with more than three or four relationships concurrently.

In summary, we decompose complex problems to temporarily reduce the number of relationships, thereby oversimplifying the problem to make it solvable utilizing a sequential logical process. We are forced to do this because our conscious reasoning facilities are easily overloaded with information.

Secondary limitations and tensions that contribute to our human problem solving difficulties include our tendency to seek a degree of accuracy that is often unrealistic and usually unnecessary. Our aversion to risk and instinctive need to survive drives us to try to predict the future with great accuracy. In this respect, we place a great deal of reliance on mathematics even though mathematical models often fail due to oversimplification of the problem situation and incorrect boundary assumptions.

We often seek to produce an optimum solution even though the problem conditions are continuously changing and, therefore, we have no benchmark that would allow us to judge whether a particular (apparent) solution is in fact optimal. In other words, under dynamic conditions there is no static benchmark available. This creates related difficulties, because our ability to interpret and judge any situation is necessarily based on comparative analysis. Subject to the experiential basis of the human cognitive system we normally have no alternative but to measure new situations with existing metrics based on past experience. However, the further the new situation deviates from past experience the more misleading the available metrics are likely to be. As a result, since we have no effective metrics for assessing new situations, we typically require a considerable period of time to correctly evaluate such situations. Accordingly, it is not unreasonable to conclude that human judgements are more influenced by the past than the present.

More comprehensively stated, the essentially experience-based nature of human cognition forces us almost always (i.e., at least initially) to apply existing methods, notions, and concepts to new situations. Therefore, our most effective problem solving capabilities utilize prototype solutions based on past experience (Gero et al. 1988). While we have become quite skilled in adapting, modifying and combining such prototype solutions, we find it very difficult to create new prototypes. As a consequence we invariably apply existing solution methods to new problem situations and develop new methods only through painful trial and error. This also leads us to generally underestimate the complexity and impact of new situations.

Human strengths

So far in this paper the discussion has centered on the apparently numerous limitations and weaknesses of human beings, particularly in respect to intellectual and emotional capabilities. Surely we human beings also have intellectual strengths. The answer is yes, of course, but with some qualifications. Certainly human learning capabilities, supported by a very large associative long-term memory, are vast. However, our rate of learning is rather slow and appears to lack efficiency. While some of this inefficiency is undoubtedly due to human communication inadequacies, the very process of progressively collecting experience by building onto existing associative knowledge structures would appear to be cumbersome and rather time consuming. It is not simply a matter of adding new knowledge elements or associating existing elements by inserting linkages, but instead patterns of neural activations (i.e., firings) have to be repeated many times before they are literally grooved into long-term memory. It is therefore not surprising that formal education takes up one quarter to one third of a human life span and involves a great deal of concentration, as well as assistance from other human beings who have acquired special teaching skills.

An important part of the human learning capability is the ability to conceptualize experiences into knowledge that we consider to be true in most cases. In this way we place emphasis on being able to deal with general conditions and consider the exceptions to the general rules to be much less important. This again exemplifies the human tendency to oversimplify a situation for the sake of being able to reach a quick solution to a problem or an explanation of an observed phenomenon. In fact, as we discover to our consternation time and again, the exceptions are often more important than the generalizations (Minsky 1990).

It must also be noted that much of human learning is involuntary and therefore virtually effortless. This applies in particular to the acquisition of low-level, largely involuntary skills such as sensory pattern matching that allows us to automatically convert data to information. For example, when we enter a restaurant we immediately recognize the furniture in the room. In fact, our eyes see only image patterns. However, these are automatically interpreted as tables and chairs by our cognitive system which has by experience related these image patterns to the appropriate symbolic entities.

At a higher level, symbolic reasoning allows us to infer knowledge from information. When our reasoning capabilities are unable to cope in complex situations that include many relationships, conceptual pattern matching (i.e., intuition) allows us to assess situations without resorting to logical reasoning. However, again there is evidence that this process is greatly facilitated by experience. Klein (1998) found that highly experienced fire captains will resort to the situation analysis methods employed by novices when they are confronted with situations outside their sphere of experience.

While the creation of new knowledge is normally the province of individuals, once such an intellectual leap has been accomplished we collectively excel in the technological exploitation of this contribution. Typically, this exploitation proceeds incrementally and involves a large number of persons, coordinated in a self-organizing fashion but willing to collaborate to leverage the capabilities of individual contributors.

Although our ability to create new knowledge is severely limited, each generation of human beings appears to be capable of producing at least a small number of individuals with far superior intellectual capabilities and insights. How these individuals create new knowledge is not at all understood, however, the ability to associate the characteristics of one situation with another apparently unrelated situation is suspected of playing an important role (Lakoff and Johnson 1980). Such analogous associations appear to provide a mental bridge that extends existing experience into unknown territory to provide a necessary basis for the creation of new knowledge through deductive reasoning.

There are other human strengths that should be briefly mentioned even though they fall outside the principal focus of this paper. Firstly, human beings appear to have some extra-sensory perception capabilities, although these have been demonstrated by very few individuals. Secondly, although human beings have generally similar needs and capabilities, they display a relatively wide range of individual differences (particularly in the intellectual and emotional realms). Thirdly, human beings tend to degrade gracefully and recover relatively quickly. For example, we recover from sleep deprivation by sleeping far fewer hours than we have previously been deprived of. Another example is the ability of our cognitive system to at least partially compensate for many forms of physical brain damage. Fourthly, human beings can be tenaciously persistent in their efforts, and unwilling to give up even under the most adverse conditions. Added to this is our ability to provide leadership and motivate others. At times we can be highly persuasive and even inspirational. And, despite all of our intellectual and emotional limitations we tend to be generally constructive and are able to maintain an innate sense of moral integrity.

However, finally, perhaps one of our greatest human strengths is the discovery of the usefulness of tools early on in our evolution. Since then we have been successful in the development and application of more and more powerful tools. Today, as discussed at the beginning of this paper we are on the verge of merging computer-based tools with the biological fabric of our very existence.

Overall assessment and conclusions

Clearly, the fundamental role that experience plays in all human endeavors cannot be overstated. Virtually all of our intellectual abilities are centered on the collection, analysis, and application of experience. While we can readily refine and extend our existing knowledge base, we have great difficulty creating new knowledge that is not founded on existing understandings. This is of course entirely consistent with the symbolic reasoning processes that are at the core of human cognitive capabilities. Reasoning cannot take place without information. The most readily available information is the experience that is held in long-term memory. New data that are collected through the senses are processed within the context of the existing experience, and may either confirm that experience or extend it in multiple ways. Sensory data that are totally inconsistent with the body of existing experience are likely to be rejected at first. Only if these inconsistencies become in themselves consistent or pose a serious emotional threat (e.g., reduced confidence) or physical threat (e.g., bodily harm) are they consciously integrated into the experience base.

As discussed previously, there are at least two consequences of these inherently human cognitive characteristics. Firstly, we are governed by a deeply rooted emotional resistance to change that allows us to abandon experience and existing understandings in only two ways. Either over extended time periods under normal conditions (i.e., in an evolutionary mode), or over short time periods under extremely threatening conditions (i.e., in a revolutionary mode). Secondly, we have great difficulties creating new knowledge and adjusting to paradigm shifts in experience and understanding. The following three questions are relevant in respect to the current homeland security situation. Are there any remedies available to ameliorate this systemic human disposition, so that it does not obfuscate or delay the adoption of more powerful information-centric computer software notions? Can we relieve the human tension that invariably accompanies periods of accelerated change, and can we implement mechanisms for facilitating the critical digital transformation from data to information and context?

It would appear, first and foremost, that there is a need for the government to assume a leadership role by defining and clearly articulating the digital transformation vision. This is of course contrary to the currently accepted position that the government should follow the lead of industry in information technology. A position based on arguments that were certainly sound in the pre-Sep.11 (2001) environment. At that time there was no compelling need to accelerate the acceptance and exploitation of intelligent software capabilities that were postulated to evolve gradually over the next two decades (Kurzweil 1999, Bell and Gray 1997). In the absence of urgency there was no need for the government to assume a major portion of the costs and risks associated with software development. Even if in the end industry passed some of the costs onto its government customers, it made sense to take advantage of competition and the cost discounting impact of the much larger public user base.

These arguments are not applicable to post-Sep.11 (2001) conditions, where forces and considerations other than commercial competition are dominant. Foremost among these is the ultimate responsibility of the government to protect its citizens from physical threats (i.e., terrorism) and economic loss. Industry, on the other hand, is driven by business objectives that translate into profit margins and marketing advantages. These objectives are more likely to be realized by small incremental advances in product capabilities rather than the revolutionary transition to an information-centric intelligent software environment. However, once this paradigm shift in computer software capabilities has been largely accomplished government can once again assume a subservient role to industry in the arena of information technology.

Secondly, there is a need for a focused and carefully orchestrated educational initiative to familiarize industry, government employees, and the public at large, with the reasons why the digital transformation is necessary and the implications of the powerful new capabilities that will become available. In parallel the government should invest in a number of relatively inexpensive short term software development projects undertaken by small leading edge companies, so that it can demonstrate to the defense industry the kinds of software systems and capabilities it is seeking. This will become particularly important in respect to the new network-centric warfare paradigm that is rapidly emerging as a core military vision in the US (Alberts et al. 1999). While the concepts of network-centric warfare reach far beyond technical considerations, information technology is nevertheless the principal enabling factor on which the successful implementation of these concepts will depend. Specifically, the tenants of network-centric warfare such as speed of command, self-synchronization and shared awareness in a geographically dispersed battlespace, require a level of decision-support that cannot be provided by a data-centric software environment.

Thirdly, there is a need for the government to persuade its own project and contract management staff to realign existing and proposed software development and acquisition plans in support of the new information-centric direction. This effort is likely to be met with a great deal of obstinate resistance, particularly in respect to acquisition programs. Most government agencies are burdened by long acquisition cycles, often in excess of four years. Even over such a prolonged period acquisition managers may still be hard pressed to keep a project on track as they guide it through the many checks and balances that government has devised with the best intentions to prevent the procurement of inferior products, price gauging, fraud, and favoritism. Naturally these acquisition managers have become adept at resisting all external forces aimed at disrupting any part of the procurement responsibility that they have been entrusted with. Yet, in the current environment of accelerated change such well intentioned obstinacy is likely to waste a great deal of time and scarce government funding.

None of these suggested approaches represents a panacea for success. All of them will be resisted by the majority of persons who are deeply ***situated*** in their experience-based environment, and who have no desire nor see the need to change. The resultant uniquely human struggle will manifest itself in many ways driven by feelings of insecurity, ill-fated attempts to dominate through control, and concurrent contradictory displays of intolerance and complacency. Much is at stake, with the outcome largely dependent on our ability to recognize and manage the distinctly ***situated*** behavioral characteristics of our human heritage.

References

Alberts D., J. Garstka and F. Stein (1999); 'Network Centric Warfare: Developing and Leveraging Information Superiority'; CCRP Publication Series, Washington, DC .

Bell G. and J. Gray (1997); 'The Revolution Yet to Happen'; in Denning and Metcalf (eds.) 'Beyond Calculation: The Next 50 Years of Computing', Springer Verlag, New York, New York (pp.5-10).

Brooks R. (1990); 'Elephants Don't Play Chess'; in Maes P. (ed.) Designing Autonomous Agents, MIT/Elsevier, Cambridge, Massachusetts (pp.3-7).

Clancey W., S. Smoliar, and M. Stefik (eds.) (1994); 'Contemplating Minds: A Forum for Artificial Intelligence'; MIT Press, Cambridge, Massachusetts.

Cooper S. I. (2002); 'Homeland Security: A View Through the Eyes of Janus'; Office of Naval Research (ONR) Workshop Series on Collaborative Decision-Support Systems'; Proceedings, September 18-19, Quantico, Virginia (pp. 127-138).

Forsyth R. (1989); 'Machine Learning: Principles and Techniques'; Chapman and Hall, New York.

Gero J., M.. Maher and W. Zhang (1988); 'Chunking Structural Design Knowledge as Prototypes'; Working Paper, The Architectural Computing Unit, Department of Architectural Science, University of Sydney, Sydney, Australia.

Gomez-Perez A. (1994); 'From Knowledge Based Systems to Knowledge Sharing Technology: Evaluation and Assessment'; Working Paper, December, Knowledge Systems Laboratory, Stanford University, Palo Alto, California.

Holland J. H. (1995); 'Hidden Order: How Adaptation Builds Complexity'; Addison-Wesley, Reading, Massachusetts.

Kauffman S. A. (1992); 'Origins of Order: Self-Organization and Selection in Evolution'; Oxford University Press, Oxford, England.

Klein, G. (1998) 'Sources of Power: How People Make Decisions', MIT Press, Cambridge, Massachusetts.

Kurzweil, R. (1999); 'The Age of the Spiritual Machines'; Viking, Penguin Group, New York.

Lakoff G. and M. Johnson (1980); 'Metaphors We Live By'; University of Chicago Press, Chicago, Illinois.

Minsky M. (1990); 'Logical vs. Analogical or Symbolic vs. Connectionist or Neat vs. Scruffy', in Artificial Intelligence at MIT., Expanding Frontiers, Patrick H. Winston (Ed.), Vol 1, MIT Press (218-243).

Pohl J. (2001); 'The Meaning of an Information-Centric Computer Environment'; InterSymp-2001, 13th International Conference on Systems Research, Informatics and Cybernetics, Baden-Baden, Germany, July 30 – August 3.

Pohl J. (1999); 'Some Notions of Complex Adaptive Systems and Their Relationship to Our World'; in Pohl J. and T. Fowler (eds.) Advances in Computer-Based and Web-Based Collaborative Systems, focus symposium: International Conference on Systems Research, Informatics and Cybernetics (InterSymp-99), Baden-Baden, Germany, August 2-6, 1999 (pp.9-24).

Pohl J., A. Chapman, K. Pohl, J. Primrose and A Wozniak (1997); 'Decision-Support Systems: Notions, Prototypes, and In-Use Applications'; Collaborative Agent Design Research Center, Technical report CADRU-11-97, Cal Poly, San Luis Obispo, CA 93407 (pp.10-11).

Taylor S. (1949); 'Science Past and Present'; Heinemann, London, England (pp.108-129).

Ushold M. and M. Gruninger (1996); 'Ontologies: Principles, Methods and Applications'; Knowledge Engineering Review, 11(2), June.

Waldrop M. (1992); 'Complexity: The Emerging Science at the Edge of Order and Chaos'; Simon and Schuster, New York.

Wood A. (2001); 'Military Experimentation: Considerations and Applications'; Office of Naval Research Workshop Series on Collaborative Decision-Support Systems'; June 5-7, Quantico, Virginia (pp. 1-8).