# Reasoning About Adversarial Intent in Asymmetric Situations

**James L. Eilbert, David M. Carmody*, Daniel Fu**, Tom Santarelli,**
**Derek Wischusen, Jason Donmoyer**

CHI Systems, Inc.
716 N. Bethlehem Pike, Suite 300
Lower Gwynedd, PA 19002
jeilbert@chiinc.com

*SYTEX, Inc.
22 Bailiwick Office Campus
Doylestown, PA 18901-2466
carmodyd@sytexinc.com

**Stottler Henke
1660 S Amphlett Blvd
San Mateo, CA 94402
fu@stottlerhenke.com

### Abstract

Counterterrorism specialists and law enforcement agencies are interested in the long-term intent or plans of the terrorists and Organized Crime members that they oppose. They often get only sporadic, incomplete, or seemingly unrelated second-hand information upon which to base their reasoning.. Some aspects of terrorist behavior are quite repetitive and regular, while terrorists go to great lengths to change and/or hide other aspects of their activity. In order to discover terrorist plans early enough to disrupt them, counter-terrorism professionals must both understand terrorist patterns of behavior and have enough evidence to begin to detect these patterns. The question addressed here is how an automated process can support plan discovery.

## Introduction

Adversarial reasoning in military contexts has traditionally focused on interactions of similar forces using similar approaches to conflict, which has come to be designated as "symmetrical warfare". Increasing attention has been paid in recent years to the problem of asymmetrical contexts, when the forces are dissimilar in composition, tactics, weapons, and approaches. Our research has focused on techniques that can be used in adversarial reasoning in asymmetrical contexts, with a focus on predicting aspects of behavior in a way that can support counterterrorism planning and operations in asymmetrical environments.

Terrorist organizations (TOs) and the counterterrorism agencies (CT) that oppose them provide examples of asymmetric forces of growing importance and the main content-domain of this project. Within this domain, the overarching presumption is that the behavior of the TOs is both focused and structured. The TO is presumed to have goals and intentions which are pursued according to an orderly process. This process is impacted or constrained not only by cultural, religious, and ethnic beliefs and values, but by the pragmatics of actions in a complex world, which requires goals to be achieved through plans involving the synchronization of actions across time, space, and often involving multiple individuals.

The missions of CT organizations in opposing TOs generally have three components -- recognition and collection of data, data analysis and hypothesis formation, and operational planning and execution. Data analysis and hypothesis formation is the focus of the Socio-Culturally Oriented Planning Environment (SCOPE) we are building for DARPA's Evidence Extraction and Link Discovery (EELD) program. The system utilizes models of terrorist activity, and information from a number of sources in order to formulate its hypotheses.

The relevant knowledge/data bases available to a SCOPE model (or a CT analyst) include:

- A set of known facts about the current mission, mainly about breaks in the terrorist organization's secrecy, and the relations among those facts;
- A catalog of TOs and general information about each of them;
- A database of terrorist cases; and
- A historical and theoretical knowledge about how terrorists organize, train, acquire financing, communicate, plan and operate, as well as , information concerning religious, ethnic, and cultural factors that may impact their operations.

Thus, SCOPE must provide mechanisms for reasoning about and combining these different sources of information.

## Overall Intent and Behavior Patterns

Generally speaking terrorists are fanatically dedicated individuals who believe they are participants in a dynamic social or political process. These people cannot, or choose not to, achieve the changes they desire through the normal political process and resort to violence. Most acts of terrorism are committed to gain publicity for their organization and purpose, to achieve political goals, or to obtain arms or financing for future operations. By performing sensational acts that attract media attention and outrage from the public, terrorists seek a government reaction that will further their cause. The Department of Defense definition of terrorism is "the calculated use of violence or the threat of violence to inculcate fear; intended to coerce or to intimidate governments or societies in the pursuit of goals that are generally political, religious, or ideological."

This definition was carefully crafted to distinguish between terrorism and other kinds of violence. The act of terrorism is defined independent of the cause that motivates it. People employ terrorist violence in the name of many causes. The tendency to label as terrorism any violent act of which we do not approve is erroneous. Terrorism is a specific kind of violence.

The official definition says that terrorism is calculated. Terrorists generally know what they are doing. Their selection of a target is planned and rational. They know the effect they seek. Terrorist violence is neither spontaneous nor random. Terrorism is intended to produce fear; by implication, that fear is engendered in someone other than the victim. In other words, terrorism is a psychological act conducted for its impact on an audience.

Finally, the definition addresses goals. Terrorism may be motivated by political, religious, or ideological objectives. In a sense, terrorist goals are always political, as extremists driven by religious or ideological beliefs usually seek political power to compel society to conform to their views. The objectives of terrorism distinguish it from other violent acts aimed at personal gain, such as criminal violence. However, the definition permits including violence by organized crime when it seeks to influence government policy. Some drug cartels and other international criminal organizations engage in political action when their activities influence governmental functioning. The essence of terrorism is the intent to induce fear in someone other than its victims, and to make a government or another audience change its political behavior.

. While the legal distinction is clear, it rarely inhibits terrorists who convince themselves that their actions are justified by a higher law. Their single-minded dedication to a goal, however poorly it may be articulated, renders legal sanctions relatively ineffective. In contrast, war is subject to rules of international law. Terrorists recognize no rules. No person, place, or object of value is immune from terrorist attack. There are no innocents.

The major objectives of TOs lead to several operating characteristics that are used to simplify the process of building SCOPE models. TOs share the following objectives:

- Exist as an Entity of Influence,
    - Organizational Structure (often using a cell model to enable continued operations if one cell is disrupted)
    - Secrecy (avoid CT detection or interdiction)
- Effect change or achieve goals through terrorist actions which requires,
    - Continuous fund raising
    - Ongoing recruitment
    - Communication and logistic plans and actions
- Attack High Value Targets
    - Mass casualties, destruction of government personnel or facilities, destruction of national symbols,
- Maximum media coverage

These objectives, in turn, produce several characteristics of TO operations, including:

Lengthy planning through execution time cycles – TO activities often unfold over long periods of time, with few explicit or overt interactions between the asymmetrical forces involved (i.e., CTs and TOs). This, plus the need for secrecy makes detailed military-style mission planning a necessity, especially when the target of the attack is difficult.

Secrecy – TOs seek to remain as invisible as possible to CTs. With fewer resources than their opponents, TOs try to keep their plans and operations entirely hidden prior to the culminating event. Likewise, CTs seek to conceal their detection means and channels from the TOs. Thus, compared to other examples of asymmetric warfare, there are relatively few direct interactions between terrorist and CT/AT groups within a mission.

Ability to truncate plans/operations – TO operations unfold in discrete steps, culminating with some overt action (e.g. assassination, bombing, kidnapping, etc.). However, indicators of CT readiness or preparation can often lead the process to be truncated. Thus, operations may or may not produce a final behavioral outcome that can be predicted, but a truncated process is to some degree an outcome in favor of the CT, while a culminating outcome is largely an outcome in favor of the TO. The tendency to abort missions as soon as the TO see indications that their activities have been detected minimizes direct confrontations. Thus, at least initially, SCOPE does not worry about C/T interactions with TOs, or how TOs behaviors or actions are impacted or altered by the actions of their opponents. This in turn allows SCOPE to treat terrorist plans as relatively static objects.

All of the characteristics listed above are strongly associated with International TOs (State Department 2000). For trans-national terrorist planning and actions greater planning, secrecy, and skill are needed since operators do not fit into their surroundings, and when more difficult targets are attacked. The line between international and local terrorism is somewhat arbitrary (Anderson and Stone 1995). In some situations, the international TO may work with a local group.

In going from conception to execution, a terrorist attack can pass through a number of phases, and the degree to which different factors influence TO behavior varies in each of these phases. There are different types of observable indicators of activity associated with each phase. For example, there are indications that when Chechen rebels found other military means cut off, they contacted established international terrorists (http://www.stratfor.com/CIS/commentary/0103162000) and began using tactics, such as plane hijacking to attempt to advance their cause.. In the process, this group moved through a number of different domains of activity. Early in the process, there were activity patterns associated with new policy formulations (e.g., the Chechen rebel decision to use terrorist tactics). We would expect cultural and organizational factors to have a strong influence in this domain. Later in the process, activity related to alliance formation took place as a top-level Chechen decision-maker made the decision to meet with leaders of an established terrorist organization. Individual psychological makeup and theories of negotiation were probably important factors in this domain of activity. Finally, the group planned and then executed the terrorist attack, i.e. hijacked an aircraft. In this domain, activity was largely dominated by the constraints of secretive military mission planning. There may have been other domains of activity that occurred, such as infrastructure building that did not lead to an immediate attack. We believe that cultural and organizational context could play a big role in detecting infrastructure building and recruitment. It is important to note that any piece of evidence collected by an agency monitoring the situation might plausibly have fit into several of these domains of activity.

The information CT analysts actually get to see is a small fraction of the observable activity. The relevant information is buried in vast amounts of noise, clutter, and deception. The analysts know that the TOs intend to do harm, however the "who, what, when where, and how" are the critical information items the TOs try to keep hidden. Another critical factor in the plan discovery type of intent inferencing the analysts do is that the earlier they discover the plan, the better the chance for disruption, interdiction, or apprehension prior to an incident.

## Patterns and Cases

The first assumption that our SCOPE system makes is that TOs will behave in a way that produces evidence that can be linked into a graph that fits a pattern. The patterns are the device that allows the system to pull graphs that are indicative of terrorist activity out of reams of other types of evidence graphs. If patterns were too tightly defined, an excessive number of them would be needed. If a pattern were too general, then a large number of incorrect matches between evidence graphs and the patterns would be found.

There are also some basic principles we would expect patterns to have.

- It should be possible to organize patterns hierarchically.
- Not all sub-patterns should be necessary to have an instance of a whole pattern
- A group of sub-patterns may have a spatial or temporal signature
- It should be possible to bias patterns (probabilities and sub-patterns) based on the context
- Patterns also need to be able to interact. A pattern of reporting can interact with a pattern of terrorist activity to create evidence. A pattern of counter-terrorist cell disruption can interact with plan terrorist activity that relies on the disrupted cell. Boundaries between interacting patterns, and sub-patterns that are part of a larger pattern are fuzzy.

The patterns that we are utilizing are based on the observation that once a target is selected mission planning concerns dominate the choice of activities. In addition, the secretive military style planning behind a TO attack is one of the most difficult aspects to change, and should provide an invariant pattern for attacks. We capture the invariance in the planning process in a hierarchy of mission planning templates (MPTs) associated with a particular domain. Currently, the MPTs are created by analysts using a form, and then translated into a set of tasks within a cognitive model. Later in the project we intend to automate the portions of the acquisition process.

The issues that must be resolved in developing an MPT hierarchy for a domain include: determining how many MPTs are needed in that domain, deciding if there are temporal, spatial, or probabilistic aspects that should be part of the MPT; and at what level of generalization or specificity should events be described within the MPT.

## The Role of Simulation

The information contained in an MPT can provide a specification for a simulator that generates plausible evidence that could be the result of real terrorist activity. There have been relatively few real terrorist missions, especially given the range of things that a TO might do. Thus, in order to apply case-based reasoning (CBR) (Stottler, Henke, & King, 1989) techniques to this problem it is necessary to augment the set of real cases with a set based on "what-if" war gaming of TO options. Thus, the MPTs provide a way of filling out the case base needed to apply CBR. In addition, simulating the MPTs can provide a sanity check on the MPTs themselves. By providing an output that is supposed to look like the real evidence that an analyst normally sees, one can ask an analyst about the quality of the simulated evidence, and use those answers to determine if there is a problem with the MPT in its current formulation.

If one can create a plausible simulator using MPTs, a reasonable question is whether we can initialize the simulator with the information we have about a current terrorist plan and get detailed information about the plan as output. We do not believe this is possible, since there would never be enough information to construct an accurate physical model of any of the situations we are interested in. The model would be so complicated that the accumulation of errors would quickly cause the model to diverge from the system it was trying to model.

A final issue related to MPTs and simulation that we want to point out relates to the difference between analyst authored MPTs, and MPTs acquired by learning techniques applied to real data. While the learning system is limited to patterns that have appeared, the analyst may also include his/her insights and observations based on their experience and expertise about how the world works for which the evidence is less clearly defined.. Thus, for the immediate future, we expect that a broader class of missions can emerge from the simulation of analyst authored MPTs than automatically acquired ones.

## The SCOPE System

The SCOPE system finds linked sets of evidence, or evidence graphs, within a large body of evidence and decides whether any of the graphs match a pattern, and thus indicate suspicious activity. Strong matches lead to hypotheses about a TO mission plan, and the system must continue to update the probabilities of these hypotheses as new data becomes available. SCOPE's basic algorithm for deducing TO mission plans from the available data is motivated by work in diagnostic expert systems. The algorithm has been designed to incorporate the steps in a manual technique currently used by intelligence analysts called the analysis of competing hypotheses (ACH).

The architecture used in a SCOPE system involves a synthesis of cognitive modeling and CBR technologies (see Figure 1). The fundamental objects that are passed between the SCOPE modules are hypotheses about the TO mission plan. One SCOPE module is based on a cognitive model of an intelligence analyst conducting situational logic (Heuer 1999), which is built using CHI System's iGEN toolset (Le Mentec, Zachary, and Iordanov 1999).

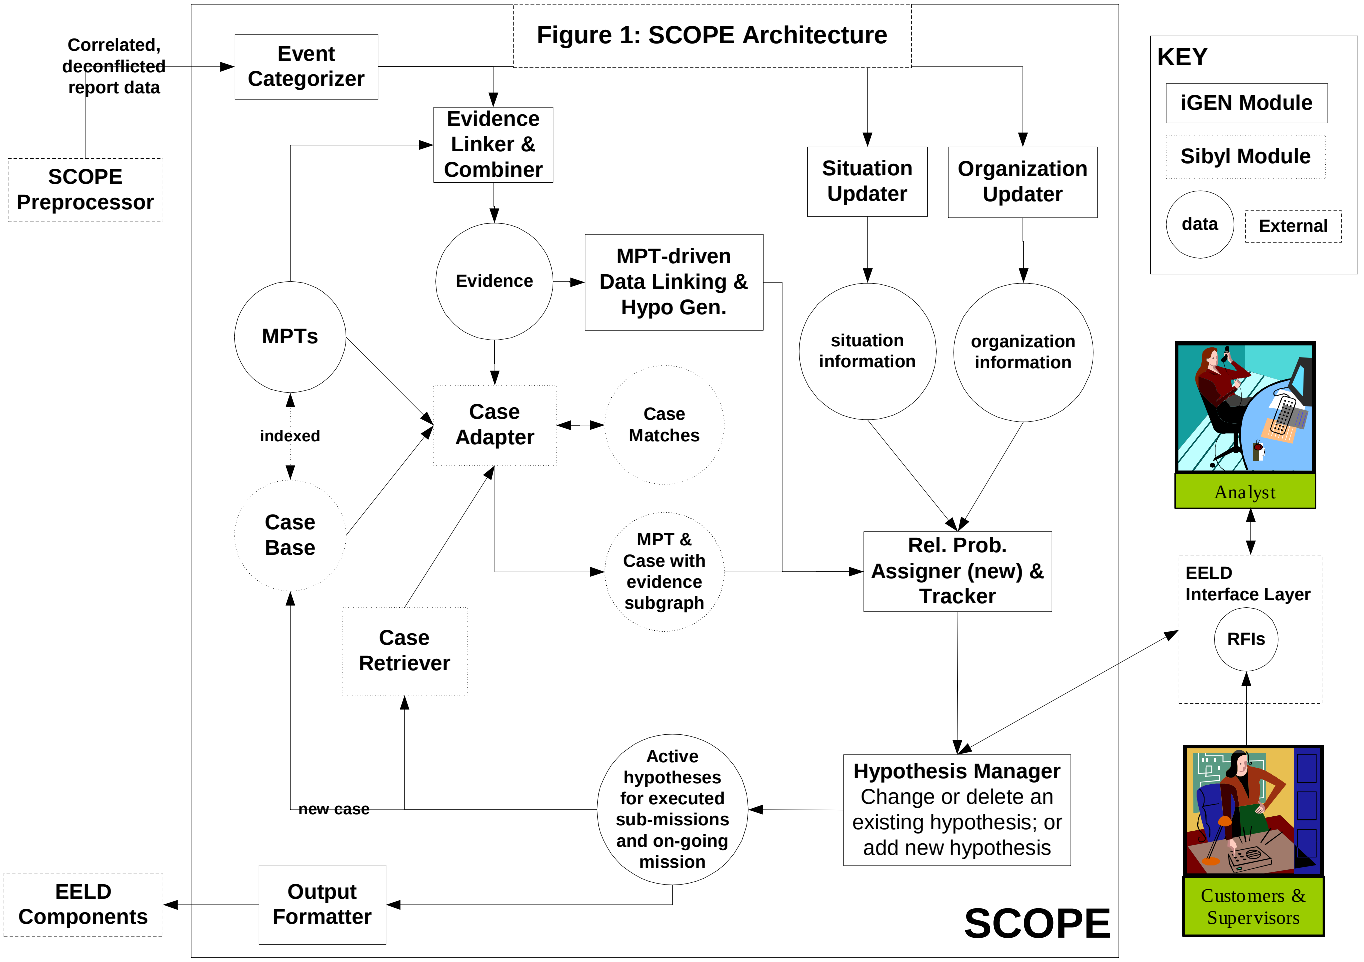
Figure 1: SCOPE Architecture

This module acts as SCOPE's primary controller. It also encodes the information in MPTs within a set of cognitive tasks, and has the meta-cognitive ability to spawn and track "what if" hypotheses about plausible mission plans. The cognitive model module reasons about how plausible hypotheses about plan components fit together, given the organizational and cultural constraints. It will also manage the active hypotheses related to MPTs taking into account the uncertainty in the evidence and sensitivity of the hypotheses.

The CBR module of SCOPE, called SIBYL, matches current evidence to plans in its case base, generating plausible hypotheses about the current TO mission plan. A case is a compact representation of:

- Primary threat events (e.g., murders)
- MPT's
- Indicators prescribed by an analyst
- Relevant ground truth

Combining and exchanging of hypotheses between SCOPE's iGEN and SIBYL modules have complementary strengths and weaknesses in generating hypotheses about mission plan execution. SIBYL needs a substantial portion of the complete evidence graph before it becomes very effective, but it is not sensitive to misconceptions an analyst may have about the TOs create mission plans. On the other hand, iGEN functions acceptably with much less complete graphs than SIBYL; however it is quite sensitive to pattern description errors that may get into an MPT. This effect was visible during our year 1 evaluation. This observation also fits nicely with Heuer's (1999) recommendation that analysts, who largely rely on situational logic, should make an effort to do more case-based reasoning.

## Situational Logic (Cognitive Model) Module

The SCOPE system was tested this year on simulated Russian Mafiya evidence primarily related to contract killing. This section describes the Situational Logic module used on that evidence, while the following section describes the initial CBR module.

iGEN stores evidence, initialization information and intermediate hypotheses on a blackboard. It stores rules and information about patterns or MPTs in a set of cognitive tasks. The evidence, which comes from the simulator as XML, is parsed and loaded into the blackboard. The final hypotheses are read off the blackboard parsed into an XML output format, and transmitted to an evaluation site.

The incoming evidence stream is processed in two stages:
Stage 1: The "sub-missions task" (or general Murder-for-hire (MFH) task) finds events that are likely to be part of a MFH pattern (the Evidence Linker and Combiner block in Figure 1) and links identified subevents to the growing MFH patterns. In a second pass, it generates assertions about missing evidence that it infers from what is still missing from the pattern, and links MFH events/sub events for use by mission identification task (the MPT-driven Data Linking and Hypothesis Generation block in Figure 1).
Stage 2: Mission task (gang war, industry takeover)
In this stage, the situational logic module identifies known gang war or industry takeover patterns and links identified MFH sub-missions. As was done in the sub-missions task, the mission task generates assertions about missing evidence inferred from what is missing in the pattern (this computation also represented by the MPT-driven Data Linking and Hypothesis Generation block in Figure 1).

Finally, the system posts high-level mission instances on blackboard for use in LD output (the Hypothesis Manager block in Figure 1).

The Situational Logic module was able to find the main mission events in 13 out of 14 simulator runs on a blind test. The different runs varied in observability, noise, and corruption.

## Case-Based Reasoning Module

The SIBYL module was tested on the same evidence that the Situational Logic module was tested, and performed at a similar level except in a few simulator runs with low observability.

The data presented to SIBYL presented special problems for standard CBR technology:

- The input is overwhelming in size, making examination of all input infeasible.
- There are no well-defined targets in the input. Targets are revealed over time, mixed together, incomplete, corrupted, embedded in noise.
- The solution size is small compared to the input.

The initial strategy for surmounting these issues was to form a "spanning case base" covering known and theoretical scenarios. Detecting TO plans in the evidence stream amounted to a search through the case memory. Thus, we match the entire case base against evidence. Hence, the CBR phase of adaptation is paramount while retrieval is secondary.

To make this approach practical, we developed methods of reducing the size of the case base, and creating fast mapping techniques. The reduction in case base size was possible by abstracting events through the Cyc ontology. Thus, we were able to condense millions of possible cases into hundreds. The creation of fast mapping techniques was achieved through domain independent heuristics computed for each individual case. The search method did not sift through the evidence; rather, it probed for certain case elements in the evidence using tightly constrained queries.

### Metrics and Visualization

We realize that the SCOPE system will not be able to do the deep reasoning about the quality of evidence that human analysts can do. However, it will be able to go through a lot more evidence per unit time than a person could do. The question is how should the SCOPE output be presented to an analyst, and how can we measure the quality of that output. The information that we believe SCOPE should visualize includes:

- Alerts about detected or suspected TO plans, operations, or combinations of indicative activities
- Data ranked by impact on current hypotheses, new hypotheses, or novelty
- Explanations of alerts or rankings
- Linked evidence graphs relating to potential missions
- Requests for Information (RFI), or detected information that matches standing Priority Information Requirements (PIRs) or Information Requirements (IRs)

The primary metric for SCOPE is issuing alerts with essentially no misses and relatively few false alarms. The accuracy of alerting was measured in the year 1 evaluation with software developed by the EELD evaluation contractor, IET. In year 2, this metric will need to be expanded to take into account whether an alert was issued with as little incremental evidence as possible.

There are additional measures of performance associated with the other types of information that SCOPE can visualize.

- Reporting all important or interesting reports (% of ranked reports that agree with reports ranked by IAs)
- Tracking all relevant hypotheses (% of hypotheses that agree with those IAs say should be tracked)
- Analysts could also supply subjective rating of the quality of explanations and RFIs.

## Next Steps

Our Year 2 goal is to achieve alert scoring similar to Year 1 on more realistic, harder evidence. A number of upgrades to the evidence that SCOPE deals with are needed before we can be confident that it can perform with the types of evidence streams that are available to analysts. The evidence stream should include more examples of early criminal or terrorist activity (rare, critical events like the murders we had in the year 1 data are often not available in this type of the early evidence.). Also, the early evidence should come in incrementally. This will add difficulty, since the evidence will not initially contain the critical, rare event in most cases that can be used to focus a search. Early data, can also lead to many competing hypotheses many of which will be wrong. If inconsistent or contradictory evidence is also allowed it will become very difficult to eliminate incorrect hypotheses. In fact, it can cause unstable interaction among a set of hypotheses.

At this point, some of the blocks in the Figure 1 architecture diagram have not been implemented. We need to develop a method for computing the impact of world events on probability assignment to hypotheses. For example, when the 9/11 terrorists are found guilty and sentenced to death,, the probability of al Qaeda activity in the following months should go up. We also need rules for when and how to make initial probability assignments. For instance, how large an evidence graph is needed before a hypothesis should be created and assigned a probability? SCOPE also needs rules for evidence combination, since pieces of evidence supporting the same hypothesis can be synergistic, mutually exclusive, or independent.

## Acknowledgements

The work described in this paper was conducted as part of DARPA’s Evidence Extraction and Link Discovery (EELD) program under Contract # F30602-01-C-0200. We would like to thank the members of the EELD community for the discussions that helped form many of the ideas presented here.

## References

1. State Department Report 2000. *Patterns of Global Terrorism 1999 – Middle East Overview*. http://www.usinfo.state.gov/topical/pol/terror/00050107.htm.
2. Heuer, R. 1999. *The Psychology of Intelligence Analysis.* Center for the Study of Intelligence, Central Intelligence Agency. http://www.cia.gov/csi/books/19104/index.html.
3. Gettys, C. 1980. *Hypothesis Generation: A Final Report on Three Years of Research, Technical Report 15-10-80.* University of Oklahoma, Decision Processes Laboratory.
4. Le Mentec, J.-C., Zachary, W., Iordanov, V. 1999. “Knowledge Tracing of Cognitive Tasks for Model-based Diagnosis.” Proceeding of 9th International Conference on Artificial Intelligence in Education, Le Mans, France.
5. Anderson, S., Stone, S. 1995. Historical Dictionary of Terrorism. Metuchen, NJ: The Scarecrow Press.
6. Stottler, R.H., Henke, A.L., King, J.A., 1989. “Rapid Retrieval Algorithms, for Case-Based Reasoning," Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI), Detroit, MI.