

Modeling and Simulating Terrorist Decision-making: A “Performance Moderator Function” Approach to Generating Virtual Opponents

Ransom Weaver³, Barry G. Silverman, PhD^{1,2}, Hogeun Shin², Rick Dubois⁴

1- Systems Engineering Dept. & Inst. For Research in Cognitive Science (IRCS)

2- Center for Human Modeling & Simulation, Computer & Info Science Dept.

3- Dept. of Asian and Middle Eastern Studies

University of Pennsylvania, Philadelphia, PA 19104-6315 (215-573-8368) barryg@seas.upenn.edu

4 - Innovative Management Concepts, Inc., Sterling, VA

ABSTRACT

An elusive goal in virtual training environments is to be able to dial up the opponent of choice – e.g., the Iraqi Republican Guard, an Hamas-type of Suicide Bomber, or the clandestine minions of Bin Laden, as a few examples. In researching alternative ways to offer such a “dial up” capability, our focus thus far is to analyze actual organizations to identify “individual differences” in the form of Performance Moderator Function scorecards and a hierarchical game theoretic approach that captures the situation, organization, population, ideologic/motivation, strategic, and tactical layers of their decision making. We are also crafting a tool that can use the scorecards to semi-automatically assemble and deploy non-traditional Semi-Automated Forces or agents on a virtual battlefield. As an initial proof of concept test, we have manually applied the approach to a scenario involving a bank bomber approaching a vehicle checkpoint. The results to date indicate the approach seems to be a useful representational formalism for generic, implementation-free models of terrorist organizations and the behavior of their members. Our next steps will be to scale up the approach and try to implement it as a terrorist generator for an existing virtual reality training environment.

1) Introduction and Overview

This paper describes a research effort to develop models of terrorist organizations that will permit us to simulate and predict what types of decisions these organizations and their agents might be likely to make. We consider a terrorist organization to be a group whose aim in using violence is primarily to achieve a psychological effect, whether on its adversary or its supporters. We assume a “rational actor” model of decision making as a point of departure and attempt to establish the utility-theoretic decision models terrorists might adopt in forming their organizations, in carrying out campaigns and operations, and in maximizing their strategic and tactical goals. The reader will recall that the rational actor model assumes only that the decisionmaker (and organization) seeks to take actions that maximize its expected utility structure –

the model places no value judgment on whether the organization’s or individual’s utility structure is warranted [1].

One task of our research is to determine how best to construct decision-theoretic models of terrorist organizations and individuals. As a working hypothesis, we believe these organization and individual decision-makers can be described via Markov Decision Processes and as repeated Bayesian games. For example, in the Maoist theory of armed struggle, the preparatory stage is characterized by actions that seek to affect separate portions of the populations of the nations or regions they are trying to influence, causing them to iterate (dynamically) through several states ranging from animosity to sympathy and membership in their movements [2]. Campaigns and missions of a given organization, also, appear to exhibit Markovian cyclicalities and draw from a reasonably finite pool of possible states and transitions. By enumerating possible states, transition probabilities, and utility levels for diverse outcomes at each new state, we are currently able to instantiate a game theoretic representation of the organizations and actors involved, as will be described. At present we have pursued the repeated games model for representing terrorist behavior in a sample scenario (Section 2), and believe this can be extended for further simulation and prediction effort. However, we are open as to which approach to pursue (e.g., Bayesian networks might prove more suited as we try to scale up) and will revisit that as the research proceeds.

Another task of this research is to cull through literature sources (news articles, web material, technical analyses, etc.) and to assemble a database that contains profiles of a reasonable sample of terrorist organizations (paramilitaries, militias, etc). This effort has already begun, and as we begin to assemble the material into a database, we hope to mine it via a variety of techniques to discover the important organizational and decision-maker profile parameters (utility structure and values), and to instantiate Bayesian prior probability estimates useful for bounding and predicting future types of decisions emanating from those organizations. Some of our initial work for assembling this database and mining


it is described in Section 3, including current utility structure illustrations (what we refer to as utility scorecards).

Lastly, we are interested in the computer generation of terrorist actors within a virtual reality world, and of the computer attempting to simulate campaigns and mission operations. So Section 4 of this paper briefly describes some of that effort as well.

2) Rational Actors and Decision Theoretic Modeling

The following diagram overviews the hierarchical nature of terrorist organization decision-making. We do not propose to describe this in any detail here, although Section 4 gives a preliminary such description. Indeed there are entire books just on a single box of this diagram (e.g., see Drake (1999) on Terrorist Target Selection) [3]. Instead we will just provide a brief discussion about how decision theoretic approaches can help us to be more precise in discussing and building models of such a process. For this discussion we shall focus on the lower three boxes primarily.

Figure 1 – Overview of the Terrorist Organization Decision Cycle



Suppose a terrorist organization exists in a world that consists of a home base and three potential targets on the other side of a military checkpoint (city hall, a bank, and a sports arena). Suppose further that its decision processes have lead the organization to decide to target the bank via a car bomb. It further knows it must get through the checkpoint to carry out the bombing operation. We can model the course of action (COA) very easily via a set of likely states of the world as shown in Figure 2 – succeeding and escaping to return home, or getting caught at the

checkpoint or bank, leading to being placed in custody or getting killed in a shootout.




Figure 2 – A Markov Chain Depicting Discrete States of the World for a Specific Terrorist Course of Action (COA)

Let us examine only the checkpoint more closely, though we could look at each node in the same way. Further, keep in mind that this example is illustrative, and no real utility values have been specified. It will be a step of the research to conduct the datamining (see next section) and to interact with experts to elicit the proper structure of the graphs such as Figure 2, and the table elements and utilities such as in Table 1.

Specifically, Table 1 shows that at the checkpoint, the terrorist could find the guard well-trained and in ready mode or in an untrained, easily surprised mode. In the latter case, the terrorist might feel there is some degree of positive utility in driving through unnoticed, shooting the guard and continuing (or dying). The only embarrassing outcome would be to get caught by an unready guard. In the second row, the terrorist has less utility for engaging in a shootout with a trained guard, as getting caught can lead to eventual release. The lower utilities in each row are for the guard and they may be similarly interpreted.

Utility Model	Drive thru Unnoticed	Shoot guard, Escape unnoticed	Shoot at guard And die	Get caught	
Guard in Unready Mode	+4	+3	+1	-1	Utility to Terrorist
	-1	-2	-2	+2	Utility to Guard
Guard in Ready Mode	+3	+1	+1	+2	Utility to Terrorist
	-2	-3	0	+3	Utility to Guard

Table 1 – Game Theory Matrix of Utilities to Each Side For Various Scenarios and Outcome Possibilities at the Checkpoint

Let us restate this more generically as a decision analysis of the course of action (COA) options (columns of the decision table). At each state, S_i , the decision analysis would enumerate the columns, COA_j , available to the decision maker

(agent) in that situation. They would be things like, "drive through unnoticed", "attack", etc. These are not intended as general options. Each situation (node) is different and would have its own COA options, although presumably these would be fairly common and would be found from datamining as described in the next section.

The agent would next assign an expected utility, u_j , to each COA. This would be based on a listing of possibility consequences, which would again be specific to the current state or situation, although ideally also tied to overall mission achievement, and generated from the datamining for Bayesian prior probabilities (and expert interviewing). In addition to a utility, each possibility outcome would also be assigned a probability, P_j , based on the agents' beliefs about achieving that possibility if the option is selected. Expected utility, E_j , is calculated in the usual way (sum across outcomes of utility times probability). This gives the agent an ability to examine strength of belief that a given COA increases mission achievement and a basis to make a decision. For example, if expected utility of the j th COA is E_j , then the rational agent attempts to maximize E as follows:

$$\text{Max } z = \text{COAScore}(i)$$

S.t.

$$\begin{aligned} \text{COAScore}_j &= E_j \\ &= \sum_{j=1}^J P_{ij} u_j \end{aligned}$$

E_j : the expected utility of COA _{j}

P_{jk} : the probability of next State k under COA _{j}

u_{jk} : the utility of next State k under COA _{j}

Of course, in a game theoretic model, the rational choice is not always the strict maximum for a single agent, but rather the maximum that can be obtained based on the opponent's actions as well. This leads to the notion of equilibrium points in the game matrix and to the idea that the agents might attempt other decision criteria other than strict maximizing [4,5]. Some alternative criteria might be:

- Minimax or even Maximin
- Decisions Under Risk
- Decisions Under Uncertainty

While this remains to be proven, we do, however, currently believe that many terrorists will tend to follow basic tenets of statistical reasoning, since they believe they are part of a campaign and that their particular COAs will be followed up by

other members [6]. Furthering their cause can be achieved even if they get caught or killed, and so they might not be inclined to adopt the Westerner's tendency to become risk averse and dominated by a non-probabilistic reasoning, such as the criterion of least regret (as applied to a soldier's life)[7].

This is not to say that terrorist reasoning is error-free, and it is likely that behavioral decision theory and other judgment biases do exist for terrorist groups. For example, group think and mob rule will often occur in crowd scenes, while terrorist organizations are known to use anchoring and adjusting from news reports about other terrorist group's actions. Similarly, the need to appease the political spectrum of a terrorist organization's supporters can also sway decision-making toward one extreme or another. And, continuing the life of the organization often becomes paramount, introducing more conservative thinking in certain respects [3]. In general we believe we can introduce such behavioral biases into our expected utility model by adjusting the utilities of a given utility structure or scorecard. Thus we can add a weighted multiplier for aggressiveness or riskiness, etc. to model such biases.

3) Database Construction and Datamining for Utility Scorecards

The information held in the database of terrorist operations is to be a compendium of attributes that can be sorted relationally for the purpose of determining what cases most closely resemble a given situation within the game environment. Each operation is a node within the database, and is composed of "scorecards" which are categories of attributes, and exist as sub-nodes.

An example of an operation entry would be:

PIRA 10/2/72 West Belfast. Attack on undercover army recon unit "MRF" killing driver of van conducting surveillance [8].

One example of many scorecards that would exist under this entry would be:


Operation environment: Urban

Where the available values in the scorecard would be:

Urban
Settled
Rural
Forest
Desert
Alpine/Arctic

The scorecard attributes are to characterize the terrorist organization, its ideology, political goals, campaign characteristics, operational environment, capabilities, tactics, and many other attributes. By means of these characterizations we hope to be able to know, when presented with a particular situation in a simulation, what a terrorist would really do. If we can know this, we can realistically bound the utility structure (COAs for a given state) and assign utilities to the actions of the terrorist agent within the simulation.

Figure 3 – Illustrative Scorecards for the Car Bomber-Checkpoint Scenario



For instance, for the terrorist car bombing operation referred to in earlier Figure 2, Figure 3 shows the terrain, with some of the attributes thereof represented as scorecards. It also shows the ideology and population support. Note that the operation takes place in an urban setting. This has a correlation with the above example of the terrorist operation in West Belfast. By assigning high probability to agent's actions that are similar to the actions recorded in the scorecards for the above PIRA operation and others that correlate highly, we can attempt to automate the generation of the game matrix (COAs, utilities) for a new state of the world.

To summarize, within the database of terrorist operations there is to be a dataset for each operation, type of organization, etc. Each datum within the dataset is a "scorecard" that records an attribute of the operation, such as the terrain, the ideology of the organization, the type of weapons used, the type of security encountered by the agent or agents, the objective of the operation, and so on. These scorecards are used to filter information into and out of the database, plus they can serve as a usage device.

Table 2 – Some of the Scorecards that have been Mined from the Database About The World In Which a Guerilla Organization Exists

Game Environment Initialization Worksheet

User specifies scorecard settings for simulation environment below:

- WA- Operational environment:**
- Urban %
 - Settled % -
 - Rural %
 - Forest %-
 - Desert %
 - Alpine/Arctic % -
 - Littoral %-
- WE- Political situation:**
- Stable democracy
 - Democracy in turmoil
 - Civil war
 - Sectarian conflict
 - Anarchy
 - Totalitarian regime
 - Ethically dominated regime
 - Foreign Military occupation

WB- Population Ratios:

- Own Group % -
- Allied group % -
- Adversary 1 % -
- Adversary 2 % -
- Neutral % -

WF- Security environment rating:

- Little security against agent
- Moderate security
- High security

WC- External sanctuary level:

- None
- Low
- Medium
- High

WG- Security environment attributes:

- ID/travel documents required
- Checkpoints/ID checks common
- Opponent has informant network
- Poor intel gathering by opponent
- Good intel gathering by opponent
- Excellent intel gathering by opponent

WD- Level of external support:

- None
- Low
- Medium
- High

Consider an example of how the database structures (scorecards) can be used as a user interface. Table 2 shows one such worksheet, which would become a program interface, by which the user enters a profile of the simulated "world" in which the terrorist organization and agents are to operate. The scorecards in Table 2 correspond to scorecards within the database of terrorist operations, and in this way the scorecards are used to screen the database for corollaries of the current situation within the simulation. We explain the screening process in the next section.

To carry this process one step further, the reader should realize there is a set of scorecards like Table 2 for each of the layers of the hierarchical decision model that was introduced as earlier Figure 1. Thus there is a set of scorecards for organization design, campaign planning, mission selection and planning, and operations (COA execution). As but one more example, Table 3 shows an illustration of the organization worksheet and some of the scorecards associated with it.

Table 3 – Some of the Scorecards that have been Mined from the Database About The Attributes of a Terrorist Organization


Organization Initialization Worksheet

- OA- Ideology:**
 - Separatism
 - Religion
 - Liberalism
 - Anarchism
 - Communism
 - Conservatism
 - Fascism
 - Single-issue
 - Organized Crime
- OB - Aims:**
 - Marxist revolution
 - Attain autonomy for ethnic group
 - Expel occupying military force
 - Enrich self/group
 - Agitate public to support authoritarian rule
 - Establish unity of religious community
 - Cast off rule of other religious group
 - Undermine authoritarian rule
 - Defend existing order
- OC - Constituency:**
 - Ethnic minority
 - Ethnic majority
 - Religious minority
 - Religious majority
 - Economic underclass
 - Economic middle
 - Economic upper class
- OD - Membership type:**
 - Intellectual/ideological
 - Ethic affiliation
 - Religious affiliation
 - Mercenary
- OE- Membership number:**
Enter number of active members
- Enter % of total population:**
 - Supporters %-
 - Potential Sympathizers %-
 - Uncommitted %-
 - Unsympathetic %-
 - Opponents %-
 - Enemies %-

Since we have defined terrorism as violence for psychological effect, it would be useful to model the opinion of the population regarding the group. To that end we have devised a model of population opinion as a series of finite states in a Markov chain, with the terrorists' course of action affecting the probability of shift from one state to another [3].

Figure 4-

Model of Population Opinion (Regarding an Organization)



This model is not one that need be developed on the virtual battlefield; rather it conceived of as

component of the case-based, offline agent generator described in the next section.

4) Putting it All Together: Virtual World Construction & Simulation Procedures

Figure 5 illustrates a decision theoretical, game theory-based approach for the modeling of a terrorist agent and organization within a computer generated simulation environment. The flowchart uses a case-based approach to establishing utility structures and weights for the agent's actions. The flowchart currently is initiated by a human filling in the preliminary scorecards, or "worksheets," that describe the situation to be modeled; however, the model could be altered for more automation and less user input. The procedure is as follows:


- 1) In The topmost box, the user characterizes the simulated "world" in which the terrorist organization and its agents are to operate. This involves filling in the weights for earlier Table 2.
- 2) In the next box, the user makes some initial characterizations of the simulated terrorist organization, and this is used to create a baseline generic terrorist organization, which is really a small set of scorecards describing the organization (as shown in earlier Table 3).
- 3) Automated Campaign Planning: The program uses these characterizations to filter the database of real terrorist operations, in order to create a terrorist campaign applicable to the present situation.
- 4) Automated Mission Selection: Selection of a target within the simulation based on the present situation, the campaign and analysis of the database for antecedents to the present conditions.
- 5) Automated mission planning: Planning the operational details of the mission based on the present situation, the target selected, and analysis of the database for antecedents to the present conditions.
- 6) Conduct Operation: Implementation of the agent, tasked to the specified mission, on the virtual battlefield and simulation to execute the COA.

In this way the program would refine the initial requirements of the user to produce a terrorist organization and agent that would behave in a realistic way within the confines of the simulation [9].

While we have emphasized machine intelligence in much of this discussion, each element in this process should have the capacity to be manually altered by the user, allowing for the steering of the semi-automated process as it accesses the case database and the various models that have been built during the run of the program. To that end we are designing a user interface that can interview the user and elicit suggested refinements

Figure 5-

Flowchart for Terrorist Decision Simulation



4) Looking forward: Present capability, JSAF Integration, and Threat Prediction

Presently we are pursuing the development of the PMF/scorecard database and its user interfaces, and we feel that this approach will be able, as output, to provide a detailed profile of a terrorist operation that is realistic for a given situation. This model terrorist will be generated offline as part of the process of implementing a terrorist agent within a simulation environment such as the military’s Joint Semi-Automated Forces or (JSAF) software environment [10]. This model generation is independent of any implementation within a simulation such as JSAF, but one that we feel could be used to provide a detailed behavioral model for a terrorist semi-automated force.

A further consideration is whether this system holds the potential for actually predicting what a particular terrorist is likely to do. It would seem that there is some scope for prediction, but that the main thrust would be simulating the operational environment. If one is attempting to simulate just a small town or region, some good predictions may be arrived at. However, if the operational scope of the terrorist is large, even international, the prospects for accurately predicting an actual act of terrorism seem small, given the vastness of potential targets in this environment.

Some other subtle aspects of modeling terrorist behavior are also problematic. The knowledge base may provide an accurate model for a military campaign for the terrorist in a given situation, but how does the campaign evolve in

reaction to countermeasures or a changing situation? Also, clandestinity is itself known to cause behavior changes such as escalating violence in the absence of central control of the operators. Also “risky shift” may occur where increasingly risky activities are undertaken in reaction to ideological and peer pressures. “Group think” may occur that suppresses rational planning and objection by minority opinions within the organization [3]. These and other factors are research dimensions we have only just begun to model. Presently our case-based model generator does not provide for such evolving aspects of terrorist behavior. Predictive modeling of terrorist behavior would seem to require their inclusion in the equation, and we hope to investigate this further.

ACKNOWLEDGEMENTS

The authors would like to acknowledge the financial support of DMSO, in general, and Colonel Forrest Crane, Dr. Ruth Willis, Dr. Phil Barry, and Lt. Col. Eileen Bjorkman, in particular. We also extend a special thanks to Martin Kleiner for his great contribution to our understanding of military systems, decision-making and planning.

REFERENCES

- [1] Savage, L.J. The Foundations of Statistics, 2nd ed., New York: Dover, 1972
- [2] Clutterbuck, R. Guerrillas and Terrorists, Chicago: Ohio University Press, 1980
- [3] Drake, C.J.M. Terrorists’ Target Selection, New York : St. Martin’s Press, 1998
- [4] von Neumann, J. & Morgenstern, O. Theories of Games and Economic Behavior, New York: Academic Press, 1964
- [5] Fudenberg D & Tirole, J. Game Theory, Cambridge, MA: MIT Press, 2000
- [6] Payne, C. & Dobson, R The Terrorists, New York: Facts On File, Inc., 1982
- [7] Wilkinson, P. “The Strategic Implications of Terrorism” www.st-and.ac.uk/academic/intrel/research/cstpv/publications1d.htm
- [8] Dillon, M. The Dirty War, New York: Routledge, 1999
- [9] Dupuy, T.N. “Military History and Case-Based Reasoning” Proceedings of a Workshop on Case-Based Reasoning, p. 125 Janet Kolodner, ed. Morgan Kaufman Publishers 1988, ISBN 0-934613-93-1
- [10] Ceranowicz, A. Nielson, P. Koss, F “Behavioral Representation In JSAF” Proceedings of the 9th Conference on Computer Generated Forces, p. 501, 2000, ISBN 1-930638-07-6