14 Methods and tools to analyze responding to, counteracting, and utilizing sociocultural behaviors[1]

Amy Sliva, Charles River Analytics Inc.

1. Introduction

The preceding chapters have explored the complexities of understanding, detecting, and forecasting sociocultural phenomena related to crucial national and international security issues. They showed how human behavior—both individual and societal—represents a confluence of cultural, political, social, economic, and historical factors interacting in a dynamic system. This chapter focuses on leveraging that information to develop effective *mitigation* policies that can respond to, counteract, or utilize these sociocultural behaviors.

In this chapter, we describe some methods and tools for mitigation analysis through computational modeling. We examine several computational approaches to course of action (COA) development, analysis, and comparison, noting the potential for positive impact on the military decision-making process (MDMP), as well as the existing technological gaps that must be filled to achieve effective sociocultural mitigation. This chapter expands on previous discussions of mitigation technologies presented in Kott and Citrenbaum (2009), which focused on the sociocultural factors and the impact on military operations. The particular applications presented here do not represent an exhaustive list of available technology solutions, but rather a sampling of computational approaches that may facilitate military decision making.

The data themselves are central to the problem of mitigation. In the modern information environment, decision makers must contend with an overwhelming quantity of data—ranging from field reports and sensor readings to communication signals and satellite imagery—which, as earlier chapters showed, may potentially contain indicators of extraordinarily intricate sociocultural phenomena. Even with the visualization, data processing, and modeling techniques already described, the amount and complexity of the information to be weighed would overwhelm any human analyst, no matter how talented or knowledgeable. In this chapter, we discuss how computational modeling can relieve some of this burden, allowing decision makers to synthesize and integrate the models, knowledge, and insights from the understand, detect, and forecast phases of analysis to provide better operational awareness. Using many methods similar to the big-data analytics employed by corporations such as Amazon or Google, analysts can exploit this wealth of sociocultural knowledge to support decision making and determine the most robust COA

[1] Approved for Public Release; Distribution Unlimited. Case Number 14-2487

in many critical situations. Of course, the vast domain knowledge and experience of military and civilian experts should not be ignored; the computational models described here are intended to augment—not automate—the decision-making process.

In Field Manual 101-5 (U.S. Army, 1997), the Department of the Army specifies a seven-step MDMP: (1) Receipt of mission, (2) Mission analysis, (3) COA development, (4) COA analysis, (5) COA comparison, (6) COA approval, and (7) Orders production. Figure 1 outlines these steps, illustrating both the commander's and the staff's roles in reaching a decision. As noted in Field Manual 101-5, decision making requires an understanding of the possible consequences of a COA and how that COA can contribute to progress from the current situation toward the stated mission goal. The MDMP is designed to facilitate comparison of multiple COAs and analyze them in light of possible enemy COAs or decisions, preventing decision makers from overlooking possible actions, outcomes, or critical aspects of the environment.

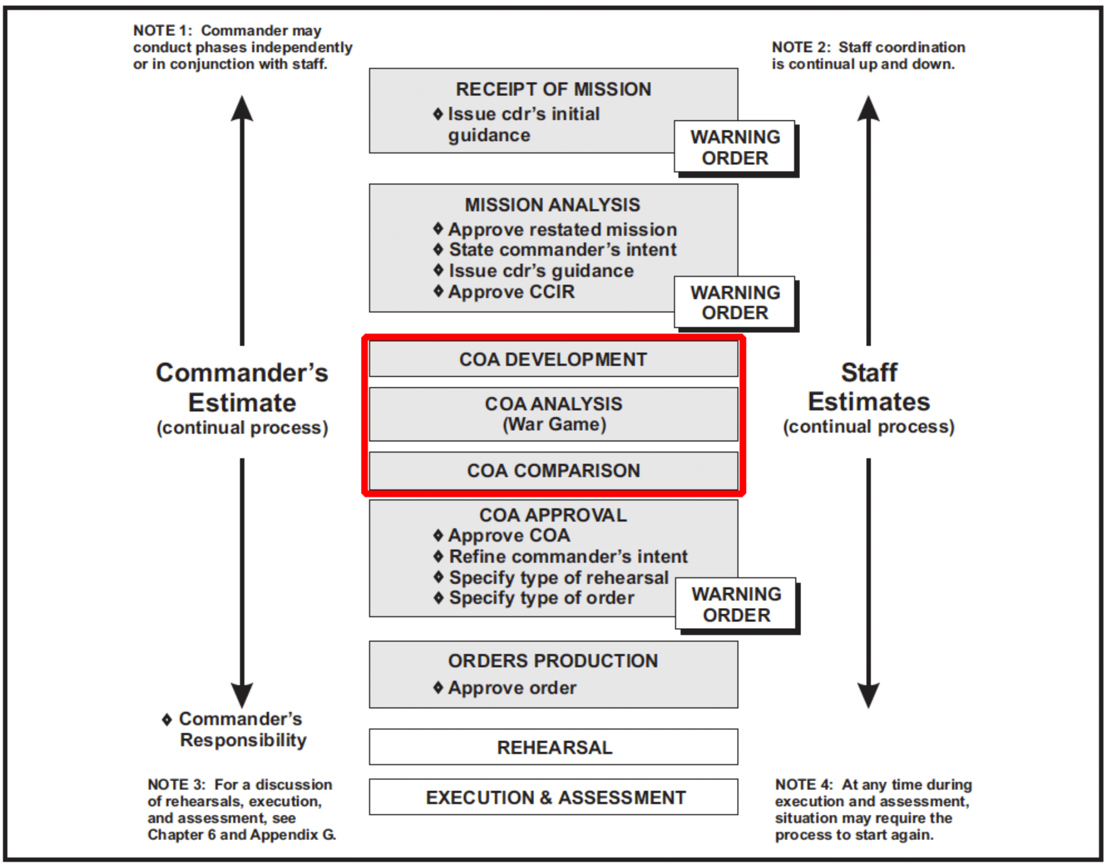

Figure 1. Military Decision-Making Process (MDMP). Adapted with permission from U.S. Army Field Manual 10-5 (1997).

Computational models can expand a commander's ability to examine possible options and account for complex behavioral dynamics when considering COAs and their short- and long-range impact.

Such models have the potential to augment this decision-making and mitigation process at several steps, outlined in red in Figure 1: COA development, COA analysis, and COA comparison.

COA development consists of producing possible COAs for further analysis and comparison (U.S. Army, 1997). All of the COAs must meet some minimal standards of suitability (i.e., adhere to constraints identified by the commander), feasibility, acceptability (in terms of the potential cost-benefit ratio), distinguishability (i.e., contains characteristics that set it apart from other COAs), and completeness. Exploratory computational modeling utilizing the behavioral models and forecasts previously developed for situation awareness can help analysts or decision makers identify COA variants that may yield better outcomes given the sociocultural dynamics of the current context. Below we describe several computational approaches to COA generation, such as abductive reasoning and artificial intelligence (AI) planning.

Typically, COA analysis involves wargaming to identify the strengths and weaknesses of a COA and major decision points that call for additional detail to account for specific situational conditions (U.S. Army, 1997). COA comparison takes this process a step further, presenting detailed analysis of which COAs are effective in particular contexts, and identifying robust options that have a high probability of success in light of the opponent's most likely (or most dangerous) behaviors (U.S. Army, 1997). Computational models capable of utilizing complex knowledge of sociocultural phenomena can facilitate both of these aspects of the MDMP. Approaches such as social network analysis and tools such as dynamic systems models and multi-agent simulations can compare the potential outcomes of various COAs (both human and model generated) under a wide range of plausible conditions to support the MDMP.

2. Mitigating Violent Behavior in Pakistan

Analysis of the sociocultural and behavioral landscape, including factors such as potential reactions of adversaries and possible roadblocks resulting from civilian cultural perspectives, is essential to the success of the MDMP. Computational models can support the MDMP by allowing officers and staff to exploit large quantities of sociocultural data and expanding their capabilities for COA analysis and comparison.

As an example, we present a complex counterterrorism scenario where a decision-making team must develop policies to mitigate the destabilizing violence perpetrated by the terror organization Lashkar-e-Taiba (LeT) in Kashmir and India. LeT has a sophisticated leadership structure and an active militant wing capable of executing complex, coordinated attacks throughout the Indian subcontinent and the world. While the organization may originally have sought only the liberation of Kashmir and establishment of an Islamist Pakistan, since the attacks on September 11th, 2001, LeT has contributed recruitment, training, personnel, and funding to the global jihadist movement. The group plays a complicated role in the sociocultural dynamics in Pakistan, using terrorist tactics to wage jihad in India, Afghanistan, and around the world, while being deeply integrated into the local social fabric through its charitable activities and provision of services that the government cannot maintain.

What counterterrorism policies might be most effective in discouraging LeT's use of violence? How can we disrupt or counter LeT propaganda conveyed via their widespread communications, education, and social services networks without threatening the security of the Pakistani citizens who rely on these services? How can the decision makers leverage their knowledge about the cultural and behavioral dynamics of LeT and the Pakistani people to achieve broader goals of security and stability within Pakistan and the region? Finding answers to these questions will require analysis of a multifaceted sociocultural landscape, augmenting human judgment and expertise with large-scale computational modeling and simulation.

3. Goal-oriented Automated COA Development

Given a particular mission and models for understanding, detecting, and forecasting sociocultural factors within the operational context, the first phase of military decision making involves developing possible COAs that could feasibly achieve the mission goals. Researchers in the field of AI have applied two main approaches to developing COAs through this type of *goal-oriented* reasoning: logical abduction and AI planning. They have used these methods, along with the predictive models discussed in Chapter 10, to construct large-scale decision support software systems used in several public and private sector decision-making processes. The computational models can generate possible mitigation policies, assess their probability of effectiveness and acceptability relative to resource and feasibility constraints, and elucidate the factors that contribute to better versus worse outcomes. Below we briefly describe the two approaches, providing some examples of possible applications and specific tools that can contribute effectively to mitigation.

3.1. Abductive Reasoning

In many everyday situations, people must interpret a set of observations, using their past knowledge of how the world works to explain the current situation. For example, a homeowner wakes up in the morning and notices that her lawn is wet. From prior experience, she knows that if it rains during the night, then her lawn will be wet in the morning. Therefore, she can conclude that it probably rained during the past night, and she may decide to carry an umbrella to work in case the bad weather continues. That, in essence, is the process of educated guessing that all people employ countless times per day: given some behavioral rule, people can form hypotheses and draw conclusions about entirely new situations. This type of reasoning is known as *abduction* (not to be confused with kidnapping!) Abductive reasoning centers on finding the most logical and relevant explanation for a set of facts. This mode of thinking is used extensively in intelligence analysis and complex decision making.

This type of reasoning does not limit people to explaining observed facts; they can use the same process to develop possible COAs instead of hypotheses (e.g., taking an umbrella to work). Basically, using abduction for COA generation involves selecting the "optimal" subset of *actions* that, when executed in the current situation, will achieve the stated goal with the highest probability, taking some possible constraints into account. The definition of optimality depends on many factors, such as the current operational environment or limitations on time or resources.

Constraints may correspond to choosing the shortest COA (either in terms of the number of distinct steps or the total time to execute the plan), the least expensive COA, etc.

The Policy Analytics Generation Engine (PAGE) develops possible COAs using both exact and approximate abductive computations, while the APT (Annotated Probabilistic Temporal) Abduction System (Simari & Subrahmanian, 2010; Molinaro, Sliva, & Subrahmanian, 2011) centers on more complex temporal aspects of COA generation. The Spatio-Cultural Abductive Reasoning Engine (SCARE) system successfully uses abduction to identify the location of weapons caches and improvised explosive devices (IEDs) in combat zones, allowing for development of specific mitigating COAs (Shakarian, Subrahmanian, & Sapino, 2009; Shakarian, Subrahmanian, & Sapino, 2011; Shakarian, Nagel, Schuetzle, & Subrahmanian, 2011). In these systems, rather than start from a set of facts or observations, users set a goal—e.g., reduce LeT use of bombings by 50%—and find an "explanation" that will enable progress from the current situation toward this goal.

If the strategic goal is to disrupt LeT's use of *fedayeen* attacks (suicide attacks) in Kashmir, the decision maker's objective is to identify ways of reducing these attacks, rather than explain how or why LeT engages in such attacks. The PAGE system utilizes contextual data about experience with combat in the region and sociocultural data regarding LeT's behavioral dynamics to determine the best way to prevent *fedayeen* attacks. The software outputs a set of actions that, when executed, would change the current situation to one in which *fedayeen* attacks have been reduced. To create deeper insight into how such COAs should be implemented, the APT Abduction System adds temporal analysis, providing the optimal time frame within which each component of the COA should be executed. This style of abductive reasoning assists analysts as they look for a strategy that can induce the desired effect, rather than perform data analysis to describe the strategy after the fact.

Abductive reasoning can also be used for more tactical-level reasoning, such as identifying and destroying enemy weapons caches in a combat zone. Rather than explain the opponent's attack strategy, the SCARE system utilizes event data and spatial sociocultural information from real combat (for example, some groups tend to place weapons caches in close proximity to IED attack sites) to indicate likely locations that should be targeted for raids or further investigation. An extension to this system is in development for locating other types of high-value targets in a counterinsurgency environment, particularly Afghanistan.

The main distinction among the abductive reasoning frameworks above results from the underlying data needed to use the systems. PAGE requires behavioral data: what events have occurred, how the group in question has responded to past strategies, etc. APT Abduction requires time-annotated behavioral data, covering not only what decisions the group has made, but also the time lags between changes in their environment, observed behaviors, and possible outcomes. SCARE relies on spatial sociocultural information: the coordinates of past attacks, the movement patterns of the group, etc. The choice of system depends on the available data and the problem a commander needs to solve. For example, if intelligence indicates that LeT will attack in India in one week, analysts might direct APT Abduction to generate all COAs that can be completed in under a week.

Event tracking, mobile phone records, satellite imagery, and the various models and data processing tools described in Chapters 1, 5, and 9 of this book can supply much of the necessary data. However, a key step in abductive reasoning and COA development relies on appropriate processing of the raw data to make them usable by the models and reasoning systems. PAGE, APT Abduction, and SCARE all rely on a data representation method known as probabilistic logic (Molinaro, Sliva, & Subrahmanian, 2011; Simari, Dickerson, Sliva, & Subrahmanian, 2012). Probabilistic (and temporal) logic has been extensively studied and used in several applications as a stochastic representation of the behavioral dynamics in international security situations, such as terrorism, drug trafficking, and ethno-political violence. Fagin, Halpern, and Megiddo (1990), Hailperin (1984), Lukasiewicz (1998), Ng and Subrahmanian (1992), and Nilsson (1986) give technical details of probabilistic logics; Mannes, Sliva, Subrahmanian, and Wilkenfeld (2008), Mannes, Michael, Pate, Sliva, Subrahmanian, and Wilkenfeld (2008), Sliva, Martinez, Simari, and Subrahmanian (2007), and Subrahmanian et al. (2007) describe sociocultural modeling applications of probabilistic logics. The benefit of this probabilistic logic-based approach results from the balance between formal reasoning—mathematical models that computers can easily manipulate—and expressive power that human users can understand.

Probabilistic logic explicitly ties behavioral and cultural dynamics to facts about the world environment, relating the probability of actions to potential contextual factors, such as the political conditions (e.g., existence of repressive government policies, occurrence of elections, corruption, etc.), sociocultural factors (e.g., religious beliefs, ethnic cleavages, leadership structures, etc.), or economic conditions (e.g., business relationships, economic grievances, macro-economic indicators, etc.). Such rules can be extracted from the behavioral or spatial data previously mentioned (Khuller, Martinez, Nau, Simari, Sliva, & Subrahmanian, 2007). Figure 2 contains a sample of a probabilistic logic representation of some behaviors of a terror organization like LeT, indicating the probability of violent actions as a result of situational factors in the environment (Subrahmanian, Mannes, Sliva, Shakarian, & Dickerson, 2012). The first line says that if the group is experiencing conflict with another rival organization (i.e., interorganizational conflict), then the probability that it will use kidnapping as a strategy lies between 0.35 and 0.45.

Probabilistic Logic Model of Terrorist Behavior

1. $\text{kidnap} : [0.35, 0.45] \leftarrow \text{interOrganizationConflicts}.$
2. $\text{kidnap} : [0.60, 0.68] \leftarrow \text{unDemocratic} \;\&\; \text{internalConflicts}.$
3. $\text{armed\_attacks} : [0.42, 0.53] \leftarrow \text{typeLeadership(strongSingle)} \;\&\; \text{orgPopularity(moderate)}.$

Figure 2. Example of probabilistic logic rules describing the behavior of a terror organization.

Unfortunately, abductive COA generation can be computationally very intensive, making it difficult to incorporate the systems mentioned above into real-time decision making. Researchers have developed a more scalable distributed framework that can approximately solve for the top *k* possible COAs (where *k* is an amount determined by the user) (Simari, Dickerson, Sliva, & Subrahmanian, 2012). This method utilizes parallel computation to process large quantities of sociocultural information simultaneously, generating several different COAs with different starting assumptions. For example, the analyst may look for COAs to fit a hypothetical situation in which Hafez Saeed, the leader of LeT, is under house arrest, or COAs that would prevent him from finding robust solutions to a variety of contingencies. Approximations can allow decision makers to develop many alternative mitigation policies in a more tractable timeframe.

Another flavor of abductive reasoning utilizes Markov logic networks (MLNs) rather than probabilistic logic. MLNs resemble probabilistic logic in that they attempt to combine the expressive power of logic (often necessary for sociocultural models) with the uncertainty computations of probability theory (Kok & Domingos, 2009; Richardson & Domingos, 2006). However, MLNs place even greater emphasis on the relational aspects of the environment, producing a network structure that can leverage algorithms, visualization tools, and human expertise developed for related graphical models such as hidden Markov models, dynamic Bayesian networks, or conditional random fields. MLNs are more suitable than those approaches for sociocultural domains because of their natural representation of complex behavioral data, such as the kind developed using the methods in Chapters 1, 5, and 9.

In MLNs, all of the environmental factors (political conditions, sources of funding, economic situations, etc.) and behaviors (kidnapping, suicide bombing, etc.) that describe a group are nodes in a graph and are connected according to their probabilistic cause and effect relationships (Pearl, 1988; Richardson & Domingos, 2006). Recall the (partial) model from Figure 2 of a terrorist group's violent behavior in the Middle East. Logic rules (1) and (2) reveal two possible relationships between the operational context and the occurrence of kidnapping. In the first case, interorganizational conflict is a potential indicator that kidnapping will occur; in the second case a combination of nondemocratic group leadership and internal conflict within the organization affects the likelihood that it will use kidnapping as a strategy. An MLN can represent these possibilities as shown in Figure 3. Because interorganizational conflict is related to kidnapping, the analyst draws a line connecting these two ideas. Since both non-democratic group leadership and internal conflict relate to kidnapping as well, the analyst connects these with the action; to indicate that these two factors must occur at the same time, the analyst also draws an edge connecting them. This small MLN compactly and visually captures these complex relationships between the environment and observed sociocultural behaviors.

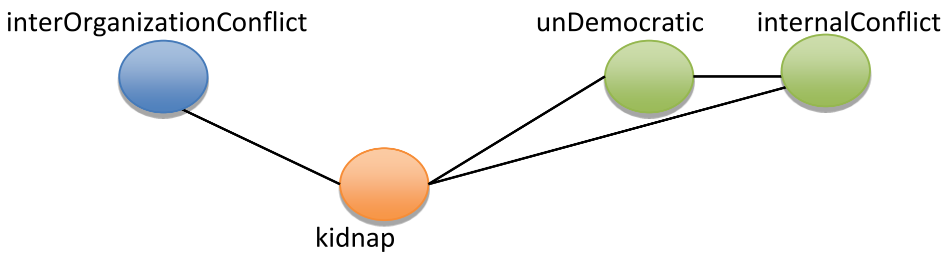

*Figure 3***.** Markov logic network representing factors that contribute to LeT's decision to use kidnapping.

Some recently developed systems can apply these MLNs to probabilistic abduction (Kate & Mooney, 2009; Singla & Mooney, 2011). These processes use the graphical structure to facilitate efficient analysis and identification of possible policies. For example, the MLN in Figure 3 clearly contains two separate sections pertaining to kidnapping. If the goal is to reduce kidnapping, then an analyst can quickly identify the relationships that lead to the lowest likelihood of kidnapping. In this case, rule (1) shows that fostering intergroup conflict might be one aspect of a COA to counter the organization's militant strategies. Because MLNs are easy to visualize, decision makers in the field can use them for real-world COA generation. Both MLNs and probabilistic logic frameworks use the concept of abductive reasoning to develop possible COAs that directly utilize the sociocultural understanding we have developed. These COAs can supplement human-generated COAs and be further compared and analyzed either manually or with additional computational models described in section 4 of this chapter.

3.2. Artificial Intelligence Planning

AI planning also explores computational models for automated COA discovery, and may provide more efficient algorithms or intuitive representations in some tactical domains. Like abduction, AI planning is goal-based, searching for possible COAs that can achieve a mission given the current operational environment and some possible constraints on time or resources.

Particularly relevant to the complex, dynamic situations confronting military decision makers is the need to plan under uncertainty. In this type of AI planning, models generate COAs in a probabilistic domain to capture uncertainty about either the outcomes of actions, the current situation, or both. In many operational environments affected by sociocultural dynamics the outcome of a COA (or part of a COA) may be uncertain—how will an adversary or civilian population react to a particular action?—but, more important, analysts may be uncertain about the exact nature of the human terrain. Their sociocultural understanding will provide information about the possible behaviors that may have an impact on a COA, but the complexity of social processes means that the resulting models may be incomplete, noisy, and uncertain. There may simply be some parts of the operational situation that military units cannot (or do not) monitor. To account for these gaps in knowledge, analysts can use a computational COA discovery model based on partially observable Markov decision processes (POMDPs) (Kaelbling, Littman, & Cassandra, 1998; Kobolov & Kobolov, 2012). POMDPs have been utilized in many real-world applications: Yost and Washburn (2000) developed methods for solving large-scale military allocation problems with POMDP algorithms, and MIT Lincoln Laboratories and the National University of Singapore are exploring a variant of POMDPs as a method for automated collision avoidance in unmanned aerial vehicles (Bai, Hsu, Kochenderfer, & Lee, 2012).

The immense complexity of the problem can make exact algorithms for solving POMDPs and generating possible COAs intractable in real-world cases. Current research in AI planning has developed several heuristic approaches that may make POMDP models more accessible to users who need to incorporate sociocultural knowledge, allowing them to access these decision-making aids in near-real time using the limited computing resources available in the field (Ross, Pineau, Paquet, & Chaib-draa, 2008). Heuristic approaches make trade-offs between accuracy and efficiency, meaning that developers must give careful consideration to the role of human decision makers in ensuring accuracy of these systems.

During COA generation it is often crucial to identify the low-level tasks that must be accomplished to achieve a broader mission goal. This type of planning can utilize an automated planning approach known as hierarchical task networks (HTNs), which explicitly decomposes the COA into smaller subtasks. For example, a decision maker may want to formulate a high-level plan to reduce the number of *fedayeen* attacks in Kashmir by LeT militants. Using an HTN system, the decision maker might decompose this problem according to the hierarchy in Figure 4. Reducing *fedayeen* attacks has two components: ending Pakistani military support of the LeT militants and disrupting the internal cohesion of the organization. The second component consists of two pieces: encouraging the LeT leadership to resign and causing the group to splinter. Of course, both of these subtasks are very broad in and of themselves, and in a real-world HTN planner they would be further divided into lower level operations to accomplish each of these goals.

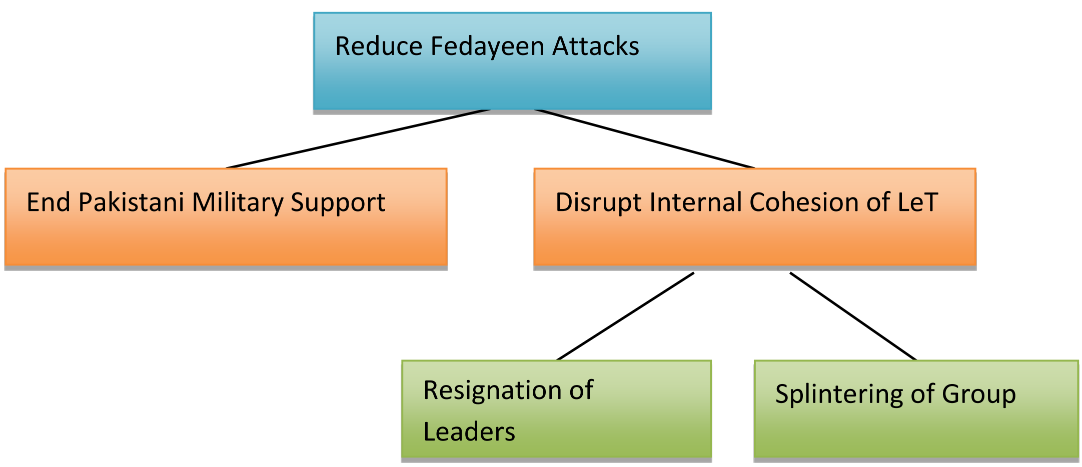

Figure 4. HTN for reducing LeT *fedayeen* attacks. The network decomposes the main goal into component parts, each of which might be further subdivided into several possible COAs.

HTNs can automatically explore different possibilities or combinations of actions to construct possible COA options that address the subtasks in the network. Because they portray the structural relationships among different aspects of a COA, planning focuses only on COAs that consist of actions that help to achieve a subtask. This directed approach leads to computation that can be more efficient for COA generation (Nau, Ilghami, Kuter, Murdock, Wu, & Yaman, 2003).

The SOCAP (System for Operations Crisis Action Planning) HTN framework, created by SRI International and tested in the field, uses a variety of data, such as threat assessments, terrain analysis, force deployment, transport capabilities, enemy behavioral factors, and operational constraints, to find possible plans. SOCAP decomposes a library of military actions into subgoals, allowing the user to select a particular objective (e.g., disrupt LeT communications), define the constraints on the environment (e.g., input the necessary sociocultural, political, geographical, etc., data), and compute possible COAs for tactical situations, such as determining equipment configurations or emergency response. The Army has also investigated HTN planning for use in small unit operations (Bienkowski, des Jardins, & Desimone, 1994; Tate, Levine, Jarvis, & Dalton, 2000). In this application SOCAP automatically translates the OPORDs (Operations Orders) produced by high levels of the military into an HTN. Users can then choose from a menu of possible objectives and use the system to generate plans for various tactical environments, such as a counterinsurgency scenario. However, few of the most recent advances in AI planning have so far been incorporated into deployed systems.

Another subset of AI planning particularly relevant to COA discovery in a complex sociocultural domain is the field of *mixed-initiative planning and scheduling* (MIPAS) (Bernstein, Beranek, Inc, & Mcdermott, 1996). MIPAS combines the efficiency and data processing capabilities of automated planners with the specialized human-in-the-loop knowledge of expert users. Where POMDP or HTN planners depend on complicated probabilistic algorithms that can be incomprehensible to actual users, decision makers may find MIPAS approaches more intuitive; rather than wait for the result of a systematic search, the user directs COA generation by dynamically indicating to the system those factors that are most relevant and must be explored or by choosing partial COAs that are most useful. As the situation on the ground changes, the user can simply click the appropriate options on the system to indicate a change in preferences. Several MIPAS planners have been developed for specific applications, such as activity plans for the Mars Rover (Bresina, Jonsson, Morris, & Rajan, 2005) and an interactive COA planning system based on the Air Force Research Laboratory's Causal Analysis Tool (CAT) (Kutur, Nau, Gossink, & Lemmur, 2004). This system allows users to iteratively create plans using CAT's causal networks for analysis and comparison. In addition, research into general MIPAS frameworks, such as the Goal Transformations (GTrans) Planner, can potentially be adapted to COA generation by using available sociocultural data as input and domain experts as the interactive users (Cox, 2003; Zhang, Cox, & Immaneni, 2002). GTrans has already been applied to emergency response situations (Immaneni & Cox, 2004).

All of the above methods can aid decision makers in identifying possible COAs in a particular operational environment. Computational decision support systems (DSSs) can enable decision makers to tap the vast sociocultural information from the understanding, detection, and forecasting phases that may be too complex for experts to analyze manually. In addition, because such computational methods are not subject to the same cognitive biases or stressful situations as human experts, these systems make it possible to automatically generate COAs that might otherwise have been overlooked.

4. Dynamic Simulations for COA Comparison and Analysis

Once analysts have developed several candidate COAs, either human or automatically generated, the next phases of the MDMP involve comparing the possible COAs and analyzing their relative strengths, weaknesses, and likelihoods of success. Computational models can facilitate comparing and analyzing COAs, allowing decision makers to consider a much broader range of possible social dynamics that could influence COA effectiveness. Computational modeling, such as dynamic simulation systems, can help identify the most robust COA, the one that will best balance resources and feasible success in a complicated sociocultural context (Bankes, 2002; 2010). The following subsections detail several of these methods, focusing on the ways that models developed for sociocultural understanding, detecting, or forecasting can be repurposed and combined for comprehensive COA analysis and comparison.

Dynamic simulation systems harness computational power to determine the most probable outcome of a COA, using the behavior of the model to analyze complex sociocultural factors and interactions that may have an impact on COA success. These simulations effectively comprise a virtual laboratory where analysts can capture the full decision-making processes of adversaries or societies of interest and evaluate how effectively a COA can mitigate such behaviors. By examining how behaviors evolve over time, they can capture emergent or adaptive behavior that may influence seemingly robust COAs. For example, a COA intended to induce internal conflict within LeT may achieve this ultimate outcome, but during the course of its application lead to increased instability in northern Pakistan. Thus, choosing the most robust COA does not depend only on the end result, but also on the intermediate effects.

4.1. Game-theoretic and Cognitive Simulation Models

Chapter 2 in this book presented several variations on rational choice and game-theoretic models. Such models can also be used to analyze possible mitigation options, allowing decision makers to explore the effects of particular actions on the decisions and behaviors of adversaries or societies (Myerson, 1991; Raiffa, 2002; Sandler & Arce, 2003). The main organizing principle for these models is rational choice theory, whereby actors make decisions to maximize potential benefits and minimize costs based on their preferences and the expected actions of other actors.

Recent work on dynamic game-theoretic models incorporates *bounded rationality*, relaxing the concept of rational behavior and including elements from cognitive or behavioral psychology (Cramerer, 2003; Gigerenzer & Selton, 2002). Simulation analyses can use such game-theoretic models in comparing the adversary behavior that will occur as a result of various potential COAs. For example, a system developed by scientists at the Duke University Institute for Homeland Security Solutions combines the adversarial aspects of game theory with the uncertainty and psychology of risk analysis, recognizing that behavior may change in probabilistic ways depending on conditions of risk (Banks, 2009). Using Bayesian statistics, the system induces a probability distribution over a possible set of outcomes at each time in the simulation, describing the possible costs/benefits of a given COA relative to a particular enemy strategy.

In a dynamic simulation, users can initialize the model with data from the current operational environment and use the probability distribution to examine all possible combinations of friendly and opponent COAs at each time point. This gives the system enough information to compute a cost-benefit analysis for each possible COA under consideration, comparing the relative probability of success in light of the most likely adversary behavior (Banks, 2009). Repeating this simulation multiple times can generate further support for a decision, enabling comparison of COAs based on the number of times the system chose each one as the optimal solution. The more frequently a COA is chosen, the more robust it may be as a mitigating policy. This risk-based framework has been used in real-world bioterrorism scenarios to analyze possible defensive COAs (Banks & Anderson, 2006).

A similar simulation approach explicitly incorporates the psychological dynamics of threat perception into game-theoretic simulations. Behavioral psychology and political science have noted the impact of threat on decision making (Huddy, Feldman, Taber, & Lahav, 2005; Pyszczynski, Solomon, & Greenberg, 2003; Schelling, 1960), where actors suddenly and drastically change their preferences under fear or stress (Kahneman & Renshon, 2006; Kahneman, Slovic, & Tversky, 1982). Researchers at Northeastern University have developed a game-theoretic simulation model based on decision making under threat (Kohentab, Pierce, & Ben-Perat, 2012; Pierce, Boulay, & Malioutov, 2011; Pierce, Boulay, & Malioutov, 2012). At every time point in the simulation, the model computes the threat perception for each friendly and enemy group based on events in the previous time; for example, if an attack occurred in that period, threat perception may have increased. The model incorporates these psychological responses to conflict events into a game-theoretic framework by determining each actor's preferences at each time point. This framework enables analysis of the efficacy of various COAs by noting the level of threat they induce throughout the simulation and the persistence or desistance of violent strategies over time.

While these approaches can enable analysis of how each possible COA will interact with an adversary's cultural and behavioral processes, it is not always possible to model truly emergent or adaptive social behavior with game-theoretic approaches because of their strict mathematical representation. The following sections describe multi-agent and complex network simulation models that can analyze COAs in light of dynamic and nonlinear behavioral models.

4.2. Complex Systems and Network Simulation Models

Complex systems science is the study of emergent behaviors within a large, interconnected system. Dynamic network simulation models offer one representation for complex systems approaches. Chapter 2 presented a detailed discussion of social network models and social influence network theory. Decision makers can use many of these same models for COA analysis to identify how reactions, perceptions, or ideas may spread through a population during implementation of a COA.

In addition, several network simulations have been developed specifically for COA analysis, utilizing structural information about sociocultural factors to map the influence of a COA. Pythia, a computational modeling tool developed by researchers at George Mason University, uses a timed influence network model to graphically represent causal relationships in a social process (Wagenhals & Levis, 2007). In this network, root nodes (nodes with no incoming edges) indicate a particular COA, including the time sequence of its component actions, while leaf nodes (nodes with no outgoing edges) represent possible final outcomes. The interior part of the network models the flow of influence and events, where a direct connection indicates a causal link, each with an associated conditional probability distribution. The system allows users to incorporate time lags into the COA events to see how the network will evolve over time.

Figure 5 traces the effects of a COA where (a) LeT's communications are interrupted, (b) LeT camps are disrupted, and (c) the Pakistani military ceases to support the group. This analysis examines the impact of the COA on the internal cohesion of LeT, which in turn has a causal impact on LeT's executing an armed attack. (Of course, in a real model the single node representing LeT internal conflicts could be a large social network representing the social dynamics of the group and its constituents.) As each element of the COA goes into effect according to the time delay, the probabilities of various aspects of LeT's internal conflicts change; this effect propagates to the likelihood of LeT's carrying out an armed attack. The Pythia simulation system makes it possible to map out the likelihood of specific outcomes—such as LeT's use of violence—as a result of social processes that evolve as a COA progresses. This facilitates comparison between various COAs or different temporal variations of the same COA (Pachowicz, Wagenhals, Pham, & Levis, 2007; Pythia: Timed Influence Net Modeler; Wagenhals & Levis, 2007).

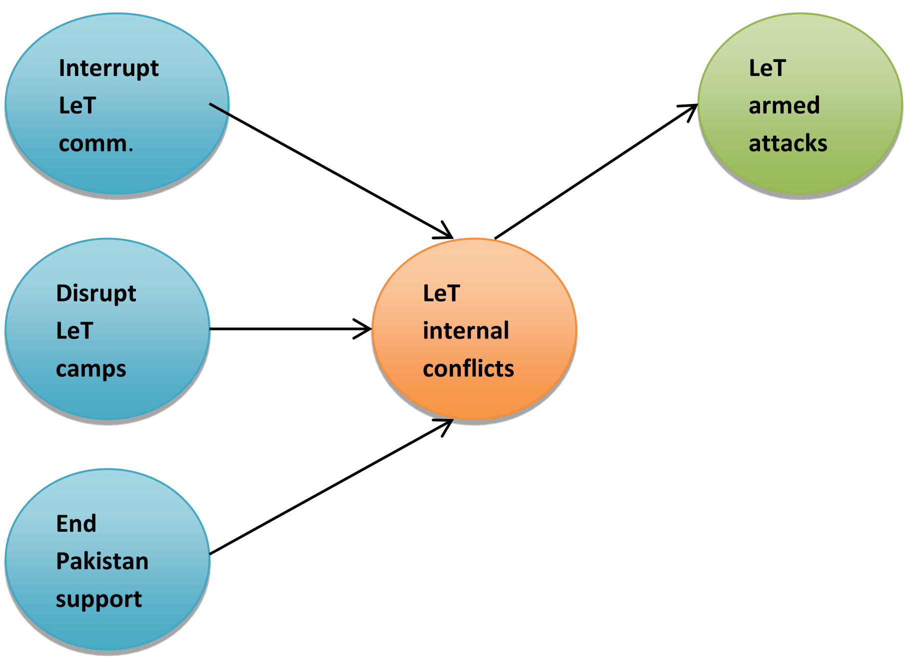

Figure 5. Influence network indicating the impact of a possible COA on the likelihood of LeT internal conflicts and behavior.

The Organization Risk Analyzer (ORA), developed by the Computational Analysis of Social and Organizational Systems lab at Carnegie Mellon University, is another dynamic network analysis and simulation tool that can facilitate COA analysis (Carley & Pfeffer, 2012; Moon, Carley, & Levis, 2008). ORA models the internal social and structural dynamics within an organization or sociocultural group. To use this tool for COA analysis, analysts must link the elements of the COA to relevant parts of a dynamic sociocultural network created for either understanding or forecasting in this domain (see Chapters 2 and 10), similar to the Pythia example in Figure 5. ORA can then analyze the impact of the COA in terms of the vulnerabilities or risks introduced into the network, how the sociocultural system changes over time, and the interaction among various different networks of interest (Carley & Pfeffer, 2012; Carley, Pfeffer, Reminga, Storrick, & Columbus, 2012; Moon, Carley, & Levis, 2008). Depending on the goals of the mission, decision makers can compare these COAs to determine the most robust option.

Basic research advances in complex science make COA analysis for mitigation a promising field for modeling (Barabasi, 2002; Bar-Yam, 2003). For example, several other models have been applied to COA analysis and comparison, such as the New England Complex Systems Institute model for comparing possible intervention policies to stem the production of poppy in Afghanistan, or support counterterrorism efforts and explore how policies with a negative perception may increase instability or volatility (Fostel & Geanakoplos, 2012; Lewling & Sieber, 2007; Widener, Bar-Yam, Gros, Metcalf, & Bar-Yam, 2012).

4.3. Multi-agent Simulations

Chapters 2 and 10 discussed the application of agent-based models to understanding and forecasting behavior in sociocultural systems. These models represent a complex system as a collection of autonomous decision-making agents that can execute behaviors according to specific sociocultural interactions (Bonabeau, 2002). The same modeling and simulation frameworks that facilitate sociocultural understanding or forecasting can also be used for COA analysis, testing the impact of potential COAs in an interactive and interconnected social environment.

One simulation tool under development for agent-based COA analysis is the Course of Action Analysis with Radio Effects Toolbox (CARET), created by RAM Laboratories (McGraw, Shao, Mumme, & Macdonald, 2009). While still in the prototype phase, CARET has provided an agent-based analysis of complex behavior in a military mission context through dynamic COA simulations that facilitate real-time decision making. CARET enables users not only to evaluate the current COA in real time, but also to simulate and compare new possible COAs to determine if a change of course might be more effective. This system is particularly designed to aid in replanning efforts: that is, updating a COA in light of new information or events in the operational environment. For example, during the execution of a policy to disrupt operations at LeT's main training compound near Muridke, Pakistan, the social situation may become unstable due to arrests of LeT members elsewhere in the country. Modelers can collect the new behavioral dynamics of this situation and incorporate them into CARET to further analyze COA viability.

Research by Laval University in Canada has also explored explicit use of Multi-Agent Geosimulation (MAGS) to support qualitative analysis of different COAs (Haddad & Moulin, 2008). This system is designed specifically to facilitate the type of "what if" analysis necessary for comparing the outcomes and side effects of various COAs in a complex sociocultural environment. Analysts can use the implanted MAGS-COA to simulate the execution of a COA in a particular virtual geographic environment that describes the geographic features of the operational theater and can change dynamically along with various sociocultural processes. Users can explore various COAs and assumptions through the agent-based simulations, modeling how the environment evolves over time as the COA is implemented. The MAGS-COA system also directly facilitates COA comparison, producing a qualitative report at the end of the simulation runs that documents the impacts of various potential COAs and their relative success in achieving particular mission goals (Haddad & Moulin, 2008). Similar analyses can also be run using the multi-agent understanding and forecasting simulations described in Chapters 2 and 10.

5. Decision Support Systems

The previous sections have briefly introduced some computational modeling methods with the potential to augment the MDMP, allowing users to leverage large quantities of sociocultural data, models, and forecasts in conjunction with their own specialized knowledge. However, many of the techniques described above are still being explored at the basic research or prototype level. To enter operational use for real-world COA discovery, analysis, and comparison, these computational models must be integrated into comprehensive applications known as DSSs. A DSS is typically defined as an interactive computer-based system intended to help decision makers use communications technologies, data, documents, knowledge and/or models to identify and solve problems, complete decision process tasks, and make decisions (Druzdzel & Flynn, 2002; Shim et al., 2002). Developers have created DSSs since the 1970s for a wide range of domains—medicine, business, air traffic control, military, etc.—that require complex decision making or logistic planning. Recent technological advances have led to a resurgence in DSS research and development, with many promising applications. The military has long utilized DSSs for non-battlefield decisions, such as personnel housing plans in the Army and the U.S. Transportation Command's Regulating and Command & Control Evacuation System (TRAC2ES), developed by Booz Allen Hamilton, for medical evacuation logistics (Booz Allen Hamilton, 2003). Several other DSS frameworks have been developed for military decision making, but have not been widely deployed.

Researchers from the University of Minnesota, the U.S. Air Force Academy, the National Aeronautics and Space Administration, the Department of the Army, and the University of Illinois have developed a suite of decision support tools for various phases of the MDMP, including generation and analysis of COAs and enemy COAs (ECOAs). FOX-GA, an AI planning-based system, uses genetic algorithms to discover COAs for military maneuvering, providing users with a range of possible options in a time-constrained environment (Schlabach, Hayes, & Goldberg, 1998). To complement these COAs, the Weasel system uses a behavioral model of an enemy or other third party to identify possible ECOAs that could result from each possible COA (Larson & Hayes, 2005; Hayes & Ravinder, 2003; Ravinder & Hayes, 2003). Weasel implements a mixed-initiative planning interface and can facilitate a "what if?" style of COA analysis. These two systems are designed to be used in conjunction with the CoRaven (or perhaps other modeling or forecasting tools) system for intelligence analysis (Jones, et al., 1998). Together, these individual DSSs can constitute a more complete decision support framework for mitigation.

The Massachusetts Institute of Technology (MIT) Lincoln Laboratory is also working on a system for integrated sensing and decision support, with specific research programs in a variety of military applications ranging from special operations to missile defense to air traffic control (Senne & Condon, 2007). This framework integrates remote sensing data, which may include sociocultural sensors such as social media or the outputs of various forecasting or detection models, into a system for decision support, simplifying the decision-making process and incorporating otherwise unavailable information.

Typically, DSSs perform most effectively in augmented decision making, facilitating human decisions rather than fully automating the analysis and final choice of COAs. These tools can help speed up the MDMP at crucial points or permit a deeper understanding of the operational environment, particularly in cases where complex sociocultural factors must be taken into consideration. This is especially true in the sociocultural domain, where even large-scale models still contain high levels of uncertainty (Bankes, Popper, & Lempert, 2006). By providing for a more efficient and complete MDMP, a DSS can help reduce the burden of complex decision making in very high-stakes situations. Of course, decision makers may be tempted to *over-rely* on a DSS; adequate training and continual assessment are necessary to ensure correct utilization of a deployed DSS.

One of the other potential benefits of using a DSS for mitigation stems from the ability of such tools to “de-bias” COA generation or analysis (Davis, Kulick, & Egner, 2005). Unaffected by the preconceptions inherent in human decision making, DSSs can promote objective decision making by generating COAs that a particular decision maker would otherwise not have considered or by analyzing possible COAs from an impartial perspective. However, some studies have shown that use of computational decision aids can actually worsen, or at the very least fail to combat, existing cognitive biases (Davis, Kulick, & Egner, 2005). In spite of the potential drawbacks of DSSs, these computational approaches to decision making can provide an excellent resource for sociocultural mitigation, making the MDMP more efficient and allowing a richer analysis of the operational environment.

6. Operational Challenges of Computational Mitigation Models

As with all new technologies, the devil is always in the details. Implementing effective DSSs or other computational models for mitigation in sociocultural environments remains a complicated technological, engineering, and training challenge.

From a research and development perspective, research must still fill many gaps before any of the above mitigation models can be deployed for operational use. Currently the MDMP can be very slow, but commanders are under extreme pressure to quickly find robust and effective COAs in severely time-constrained environments. While the computational models described above allow users to consider far more data regarding sociocultural dynamics than a human alone could analyze, the high complexity of the data means that these models may not necessarily arrive at a solution faster.

In many cases, these computational models and simulations not only require excessive running time, but also consume massive amounts of computing resources to process the large quantity of data and complexity of the sociocultural dynamics. While high-performance and distributed computing platforms, such as Apache Hadoop (Apache, 2012), are becoming increasingly available off the shelf, they still require substantial computation resources, usually available only in garrison or at the largest deployed facilities. Because DSSs are intended to augment expert decision making, the COA discovery or comparison should not be performed remotely by users who lack first-hand knowledge of the situation on the ground. The true challenge, therefore, is to provide decision support in real time so that decision makers can use these mitigation models in the theater of operations. To enable such use in the field, researchers must develop large-scale systems that do not depend on physical access to high-performance computing systems.

The previous discussion also assumes that large-scale DSSs can be developed from various component sociocultural models and that information from the models for understanding, detecting, or forecasting behavior can be directly transferred to one of the mitigation models. However, given that many models utilize unique data representation and output formats, creating a complete system can represent a serious undertaking. Developers must conduct further research in fields such as multi-modeling, model interoperability, and data fusion to generate models capable of manipulating the outputs of sociocultural understanding or predictive models into proper inputs for mitigation simulations in a standardized format.

Further, it is crucial to develop models transparent enough to foster trust by real-world users, without requiring them to become experts in computational modeling. A simulation system so complex that only an expert in computer science or mathematics can run it will not be effective for decision support of mitigation policies. Transparency can be facilitated through adequate training on how to apply these models and interpret the results when identifying possible COAs and analyzing their implementation in a particular sociocultural environment. However, significant research in human computer interaction is still necessary to determine how to design user interfaces for these computational mitigation models to ensure effective, intuitive use, and determine how best to incorporate them into the MDMP.

The final problem that affects all computational models described in this book centers on data. While many of these reasoning and simulation approaches are designed to handle large quantities of data, the outcomes will only be useful for discovering, analyzing, and comparing COAs if the input data (or the data used for the underlying models for understanding, detecting, and forecasting sociocultural dynamics) are reliable. For many social analyses, Internet social media present a new stream of human-generated data that may serve as an open source “sensor” providing insight into social processes. However, other tasks require data extraction from widely available, unstructured textual sources (news media, public statements, etc.) or qualitative information (interviews, surveys, etc.) collected through scholarly fieldwork or official human intelligence channels. As Chapter 13 discussed in depth, integrating this wealth of information into computational mitigation models presents a major challenge.

7. Summary

This chapter has examined the potential for computational models to aid decision makers in developing, analyzing, and comparing possible COAs for mitigating sociocultural factors in the environment. Such tools have the potential to significantly ease the analytic and decision-making burden faced by humans in the field, allowing users to process large quantities of complex social or behavioral data that would otherwise be inaccessible. Using large-scale data analytics similar to those applied by companies such as Amazon or Google, these modeling frameworks can leverage the wealth of sociocultural information gained from understanding, detecting, and forecasting aspects of the human terrain. Mitigation models, therefore, can augment the existing MDMP, facilitating the development of robust COAs while also exploiting the expertise and knowledge of decision makers in the field. However, many of these enabling technologies are still in the early stages of research and development, and system developers must overcome significant technological and training challenges before mitigation models will become widely operational for real-world situations.

References

Apache (2012). *Hadoop*. Retrieved from http://hadoop.apache.org

Axelrod, R. (1984). *The evolution of cooperation.* New York, NY: Basic Books.

Bai, H., Hsu, D., Kochenderfer, M. J., & Lee, W. S. (2012). Unmanned aircraft collision avoidance using continuous-state POMDPs. *Robotics: Science and Systems VII*, 1.

Bankes, S. (2002). Tools and techniques for developing policies for complex and uncertain systems. In *Proceedings of the National Academy of Sciences, 99* (Suppl 3), 7263–7266.

Bankes, S. (2010, November). Robustness, adaptivity, and resiliency analysis. In *Complex Adaptive Systems –Resilience, Robustness, and Evolvability: Papers from the AAAI Fall Symposium, v.FS-10-03* (pp 2-7). Menlo Park, CA: AAAI Press.

Bankes, S., Popper, S., & Lempert, R. (2006, March). Supporting decisions with (less than perfect) social science model. *IEEE Aerospace Conference.* Big Sky, MT.

Banks, D. (2009). *Adversarial risk analysis: Decision-making when there is uncertainty during conflict.* Durham, NC: Duke University, Institute for Homeland Security Solutions,

Banks, D., & Anderson, S. (2006). Game theory and risk analysis in the context of the smallpox threat. In A. Wilson, G. Wilson, & D. Olwell (Eds.), *Statistical methods in counterterrorism* (pp. 9–22). New York, NY: Springer.

Barabasi, A.-L. (2002). *Linked: The new science of networks.* Cambridge, MA: Perseus.

Barry, P. S., & Koehler, M. T. (2005, December). Leveraging agent-based simulation for rapid course of action development. In M. E. Kuhl, N. M. Steiger, F. B. Armstrong, & J. A. Joines (Eds.), *Winter Simulation Conference.* Orlando, FL.

Bar-Yam, Y. (2003). *Dynamics of complex systems.* Boulder, CO: Westview Press.

Bernstein, M. H., Beranek, B., Inc., N. & Mcdermott, D. V. (1996). Issues in the development of human-computer mixed-initiative planning. In B. Gorayska & J. L. Mey (Eds.), *Cognitive technology: In search of a humane interface* (pp. 285–303). Amsterdam, Netherlands: Elsever Science B.V.

Bienkowski, M. A., des Jardins, M. E., & Desimone, R. V. (1994, February). SOCAP: System for operations crisis action planning. In *Proceedings of the ARPA/Rome Laboratory Planning Initiative Workshop.* Arlington, VA: SRI International.

Bonabeau, E. (2002, May 14). Agent-based modeling: Methods and techniques for simulating human systems. *Proceedings of the National Academy of the Sciences, 99*(3), 7280-7287.

Booz Allen Hamilton. (2003). TRANSCOM regulating and command and control evacuation system (*TRAC2ES).* Retrived from http://www.dote.osd.mil/pub/reports/FY2000/airforce/00trac2es.html

Bresina, J., Jonsson, A., Morris, P., & Rajan, K. (2005, August). Mixed-initiative activity planning for Mars rover. In *Proceedings of the 19th International Joint Conference on Artificial Intelligence* (pp. 1709-1710). San Francisco: Morgan Kaufman Publishers.

Carley, K. M., & Pfeffer, J. (2012, July). *Dynamic network analysis (DNA) and ORA*. Paper presented at 2nd International Conference on Cross-Cultural Decision Making. San Francisco, CA.

Carley, K. M., Morgan, G., Lanham, M., & Pfeffer, J. (2012, July). *Multi-modeling and sociocultural complexity: Reuse and Validation*. Paper presented at 2nd International Conference on Cross-Cultural Decision Making. San Francisco, CA.

Carley, K. M., Pfeffer, J., Reminga, J., Storrick, J., & Columbus, D. (2012). *ORA user's guide 2012.* Pittsburgh, PA: Carnegie Mellon University, Center for the Computational Analysis of Social and Organization Systems.

Cox, M. T. (2003, August). *Planning as mixed-initiative goal manipulation*. Paper presented at the Workshop on Mixed-Initiative Intelligence Systems at 18th International Conference on Artificial Intelligence. Menlo Park, CA: AAAI Press.

Cramerer, C. (2003). *Behavioral game theory.* Princeton, NJ: Princeton University Press.

Davis, P. K., Kulick, J., & Egner, M. (2005). *Implications of modern decision science for military decision-support systems.* Santa Monica, CA: The RAND Corporation.

Druzdzel, M. J., & Flynn, R. R. (2002). Decision support systems. In A. Kent (Ed.), *Encyclopedia of library and information science.* New York, NY: Marcel Dekker, Inc.

Fagin, R., Halpern, J. Y., & Megiddo, N. (1990). A logic for reasoning about probabilities. *Information and Computation, 87*(1–2), 78–128.

Fostel, A., & Geanakoplos, J. (2012). Why does bad news increase volatility and decrease leverage? *Journal of Economic Theory, 147*(2), 501-525.

Gigerenzer, G., & Selton, R. (2002). *Bounded rationality.* Cambridge, MA: MIT Press.

Haddad, H., & Moulin, B. (2008, March). Multi-agent geosimulation in support to "what if" courses of action analysis. In *Simutools '08. Proceedings of the 1st International Conference On Simulation Tools and Techniques For Communications, Networks and Systems & Workshops* (Article No. 65).

Hailperin, T. (1984). Probability logic. *Notre Dame Journal of Formal Logic, 25*(3), 198–212.

Hayes, C. C., & Ravinder, U. (2003). Weasel: A user guided enemy course of action generator. *Lecture Notes in Artificial Intelligence (Subseries of Lecture Notes in Computer Science)*, *v 2774 PART 2,* 692-696.

Huddy, L., Feldman, S., Taber, C., & Lahav, G. (2005, July). Threat anxiety and support for antiterrorism policies. *American Journal of Political Science, 49*(3), 593-608.

Immaneni, T., & Cox, M. T. (2004, January). GTrans: An application for managing mixed-initiative collaborative planning during emergency response situations. In *Proceedings of the 2004 International Symposium on Collaborative Technologies and Systems: CTS 2004* (pp. 121-126).

Jones, P. M., Hayes, C. C., Wilkins, C., Bargar, R., Sniezek, J., Asaro, P., . . . Schalbach, M. J. (1998, October). CoRAVEN: Modeling and design of a multimedia intelligent infrastructure for collaborative intelligence analysis. In *Proceedings of the International Conference on Systems, Man, and Cybernetics* (pp. 914–919). IEEE.

Kaelbling, L. P., Littman, M. L., & Cassandra, A. R. (1998). Planning and acting in partially observable stochastic domains. *Artificial Intelligence, 101*, 99–134.

Kahneman, D. J., & Renshon, J. (2007). Why hawks win. *Foreign Policy*, 34-38. Retrieved from http://www.foreignpolicy.com/articles/2006/12/27/why_hawks_win#sthash.AHRzbny1.dpbs

Kahneman, D. J., Slovic, P., & Tversky, A. (Eds.) (1982). *Judgment under uncertainty: Heuristics and biases.* Cambridge, UK: Cambridge University Press.

Kakas, A. C., Kowalski, R. A., & Toni, F. (1993). Abductive logic programming. *Journal of Logic and Computation, 2*(6), 719–770.

Karsai, G., Maroti, M., Ledeczi, A., Gray, J., & Sztipanovits, J. (2004). Composition and cloning in modeling and meta-modeling. *IEEE Transactions on Control Systems Technology, 12*(2), 263-278.

Kate, R. J., & Mooney, R. J. (2009, July). *Probabilistic abduction using Markov logic networks*. Paper presented at the IJCAI-09 Workshop on Plan, Activity, and Intent Recognition (PARI-09). Pasadena, CA.

Khuller, S., Martinez, V., Nau, D., Simari, G. I., Sliva, A., & Subrahmanian, V. (2007). Computing most probable worlds of action probabilistic logic programs: Scalable estimation for 10^30,000 worlds. *Annals of Mathematics and Artificial Intelligence, 51*, 295-331.

Kobolov, M., & Kobolov, A. (2012). Planning with Markov decision processes: An AI perspective. In R. J. Brachman, W. W. Cohen, & T. Dietterich (Eds.), *Synthesis lectures on artificial intelligence and machine learning.* San Rafael, CA: Morgan & Claypool.

Kohentab, S., Pierce, G., & Ben-Perat, G. (2012). *Decision making under threat: Israeli and Palestinian public opinion*. (Unpublished doctoral dissertation). Northeastern University, Boston, MA. Retrieved from http://hdl.handle.net/2047/d20002703.

Kok, S., & Domingos, P. (2009, June). *Learning Markov logic network structure via hypergraph lifting*. Paper presented at the 26th Annual International Conference on Machine Learning. Montreal, Quebec, Canada: ACM.

Kott, A., & Citrenbaum, G. (Eds.). (2010). *Estimating imapct: A handbook of computational methods and models for anticipating economic, social, political and security effects in international interventions.* New York, NY: Springer.

Kutur, U., Nau, D., Gossink, D., & Lemmur, J. F. (2004, October). *Interactive course-of-action planning using causal models.* Paper presented at the Third International Conference on Knowledge Systems for Coalition Operations *(KSCO-2004).* Pensacola, FL.

Larson, A. D., & Hayes, C. C. (2005, September). An assessment of Weasel: A decision support system to assist in military planning. In Proceedings from the *Human Factors and Ergonomics Society, 49*(3), 287–291.

Levis, A. (2012). *Multi-modeling and meta-modeling of adversaries and coalition partners.* Fairfax, VA: George Mason University, Electical and Computer Engineering.

Levis, A. H., Wagenhals, L. W., & Zaidi, A. K. (2012, June). Multi-modeling of adversary behaviors. In *Proceedings from the International Conference on Intelligence and Security Informatics* (pp. 185–189). Arlington, VA: IEEE.

Lewling, T., & Sieber, O. (2007, January). Using systems dynamics to explore effects of counterterrorism policy. In Proceedings of the *40th Annual Hawaii International Conference on Systems Sciences.* Honolulu, HI: IEEE.

Lukasiewicz, T. (1998, August). Probabilistic logic programming. In *Proceedings of the 13th European Conference on Artificial Intelligence.* Brighton, UK.

Mannes, A., Michael, M., Pate, A., Sliva, A., Subrahmanian, V., & Wilkenfeld, J. (2008, April). Stochastic opponent modeling agents: A case study with Hezbollah. In *Proceedings of the International Conference on Social Computing, Behavioral Modeling, and Prediction.* Phoenix, AZ.

Mannes, A., Shakarian, J., Sliva, A., & Subrahmanian, V. (2011, September). A computationally-enabled analysis of Lashkar-e-Taiba attacks in Jammu & Kashmir. In *Proceedings of the European Intelligence and Security Informatics Conference* (pp. 224–229). Athens, Greece: Institute of Electrical and Electronics Engineers (IEEE).

Mannes, A., Sliva, A., Subrahmanian, V., & Wilkenfeld, J. (2008, September). Stochastic opponent modeling agents: A case study with Hamas. In *Proceedings of the 2nd International Conference on Computational Cultural Dynamics (ICCCD).* College Park, MD: University of Maryland.

McGraw, R. M., Shao, G., Mumme, D., & Macdonald, R. A. (2009). Design of an agent-based course of action (COA) analysis with radio effects toolbox. *International Journal of Intelligence Control and Systems, 14*(1), 104–114.

McIlraith, S. A. (1998). *Logic-based abductive inference.* Stanford University, CA. Knowledge Systems Laboratory.

Molinaro, C., Sliva, A., & Subrahmanian, V. (2011, May). Abduction in annotated probabilistic temporal logic. In J.P. Gallagher, & M. Gelfond (Eds.), *International Conference on Logic Programming (Technical Communications). (pp. 240-250)*. Lexington, KY.

Moon, I. C., Carley, K. M., & Levis, A. H. (2008, April). Vulnerability assessment on adversarial organization: Unifying command and control structure analysis and social network analysis. Proceedings of the: *SIAM International Conference on Data Mining Workshop on Link Analysis, Counterterrorism and Security.* Atlanta, GA.

Myerson, R. (1991). *Game theory: Analysis of conflict.* Cambridge, MA: Harvard University Press.

National Counterterrorism Center. (n.d.). *Worldwide Incident Tracking System (WITS)*. Retrieved 2012, from https://wits.nctc.gov

Nau, D., Ilghami, O., Kuter, U., Murdock, W. J., Wu, D., & Yaman, F. (2003). SHOP2: An HTN planning system. *Journal of Artificial Intelligence Research, 20*, 379-404.

Neema, H., Hemingway, G., Sztipanovits, J., & Karsai, G. (In Press). The C2 wind tunnel: Rapid synthesis of HLA-based heterogeneous simulations via a model-based integration approach. *Simulation*.

Neema, S., Bapty, T., Koutsoukos, X., Neema, H., Sztipanovits, J., & Karsai, G. (2009, July). *Model based integration and experimentation of information fusion and C2 systems*. Paper presented at 12th International Conference on Information Fusion*.* Seattle, WA: IEEE.

Ng, R., & Subrahmanian, V. (1992). Probabilistic logic programming. *Information and Computation, 101*(2), 150-201.

Nilsson, N. J. (1986). Probabilistic logic. *Artificial Intelligence, 28*, 71–87.

Oren, T. I. (1987, July). Model update: A model specification formalism with a generalized view of discontinuity. In *Proceedings of the 1987 Summer Computer Simulation Conference* (pp. 698–694). Montreal, Canada: SCS Press.

Oren, T. I. (1991). Dynamic templates and semantic rules for simulation advisors and certifiers. In A. P. Fishwick, & B. R. Modjeski (Eds.), *Knowledge-based simulation: Methodology and application* (pp. 53-76). New York, NY: Springer-Verlag.

Pachowicz, P. W., Wagenhals, L. W., Pham, J., & Levis, A. H. (2007, April). Building and analyzing timed influence net models with internet enabled Pythia. In K. Schum, & D. A. Trevisani (Eds.), *Modeling and Simulation for Military Operations II; Proceedings of the SPIE.* Bellingham, WA: SPIE.

Paul, G. (1993, April). Approaches to abductive reasoning: An overview. *Artificial Intelligence Review, 7*(2), 109–152.

Pearl, J. (1988). *Probabilistic reasoning in intelligent systems: Networks of plausible inference.* San Francisco, CA: Morgan Kaufmann.

Pierce, G., Boulay, C., & Malioutov, M. (2011, July). *An inter-group conflict model integrating perceived threat, vested interests, and alternative strategies for cooperation*. Paper presented at SIAM Conference on Control and Its Applications*.* Baltimore, MD.

Pierce, G., Boulay, C., & Malioutov, M. (2012). An inter-group conflict model integrating perceptions of threat and vested interest: Extending rational choice to incorporate psychological dynamics. *Game Theory and Applications, 16*.

Pious, S. (1993). *The psychology of judgment and decision making.* New York, NY: McGraw-Hill, Inc.

Putnam, R. D. (1988). Diplomacy and domestic politics: The logic of two-level games. *International Organization, 42*(3), 427–460.

Pyszczynski, T., Solomon, S., & Greenberg, J. (2003). *In the wake of 9/11: The psychology of terror*. Washington, D.C.: American Psychological Association*.*

Pythia: Timed Influence Net Modeler. (n.d.). Retrieved from SAL-GMU: http://sysarch.gmu.edu/main/software

Raiffa, H. (2002). *Negotiation analysis.* Cambridge, UK: Cambridge University Press.

Ravinder, U., & Hayes, C. C. (2003, October). Weasel: An automated planner that users can guide. In *Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics* (pp. 953-955). Washington, DC: IEEE.

Richardson, M., & Domingos, P. (2006). Markov logic networks. *Machine Learning, 62*(1–2), 107–136.

Riedel, S., & Meza-Ruiz, I. (2008, August). Collective semantic role labelling with Markov logic. In *Proceedings of the Twelfth Conference on Computational Natural Language Processing* (pp. 193-197)*.* Manchester, UK.

Ross, S., Pineau, J., Paquet, S., & Chaib-draa, B. (2008). Online planning algorithms for POMDPs. *Journal of Artificial Intelligence Research, 32*, 663–704.

Sandler, T., & Arce, M. (2003). Terrorism and game theory. *Simualtion and Gaming, 34*, 319–337.

Schelling, T. C. (1960). *The strategy of conflict.* Oxford, UK: Oxford University Press.

Schindler, J. V., & Mraz, T. (2008, February). Agent-based modeling for real-time decision-support for point-of-distribution managers during influenza mass vaccination. *AMIA Annual Symposium 2008,* 1124*.* Atlanta, GA.

Schlabach, J. L., Hayes, C. C., & Goldberg, D. E. (1998). FOX-GA: A genetic algorithm for generating and analyzing battlefield courses of action. *Evolutionary Computation, 7*(1), 45-68.

Senne, K. D., & Condon, G. R. (2007). Integrated sensing and decision support. *Lincoln Laboratory Journal, 16*(2), 237-243.

Shakarian, P., Nagel, M. K., Schuetzle, B. E., & Subrahmanian, V. (2011, August). Abductive inference for combat: Using SCARE-S2 to find high-value targets in Afghanistan. In *Proceedings of the 25th AAAI Conference on Artificial Intelligence and the 23rd Innovative Applications of Artificial Intelligence Conference (AAAI-11/IAAI-11)* (pp. 1689-1694)*.* San Francisco, CA.

Shakarian, P., Subrahmanian, V., & Sapino, M. L. (2009, December). *SCARE: A case study with Baghdad*. Paper presented at the 3rd International Conference on Computational Cultural Dynamics (ICCCD-09). College Park, MD.

Shakarian, P., Subrahmanian, V., & Sapino, M. L. (2011). GAPs: Geospatial abduction problems. *ACM Transactions on Intelligent Systems and Technology, 3*(1), 7.

Shim, J. P., Warkentin, M., Courtney, J. F., Power, D. J., Sharda, R., & Carlsson, C. (2002). Past, present, and future of decision support technology. *Decision Support Systems, 33*(2), 111–126.

Simari, G. I., & Subrahmanian, V. (2010, July). Abductive inference in probabilistic logic programs. In *Technical Communications of the 26th International Conference on Logic Programming, ICLP 2010* (pp.192-201). Florence, Italy.

Simari, G. I., Dickerson, J., Sliva, A., & Subrahmanian, V. (2012). Parallel abductive query answering in probabilistic logic programs. *ACM Transactions on Computational Logic, 14*(2), 12.

Simon, H. A. (1987). Bounded rationality. In J. Eatwell, M. Millgate, & P. Newman (Eds.), *The new palgrave: A dictionary of economics.* London, UK: Macmillan.

Singla, P., & Mooney, R. J. (2011, August). Abductive Markov logic for plan recognition. In *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence*, 1069–1075. San Francisco, CA.

Sivillo, J. K., Ahlquist, J. E., & Toth, Z. (1997, December). An ensemble forecasting primer. *Weather and Forecasting, 12*, 809–818.

Sliva, A., Martinez, M. V., Simari, G. I., & Subrahmanian, V. (2007, August). SOMA models of the behaviors of stakeholders in the Afghan drug economy: A preliminary report. Proceedings of the *International Conference on Computational Cultural Dynamics (ICCCD).* College Park, MD: University of Maryland. ACM Press.

Subrahmanian, V. S., Albanese, M., Martinez, V., Reforgiato, D., Simari, G. I., Sliva, A., . . . Wilkenfeld, J. (2007). CARA: A cultural reasoning architecture. *IEEE Intelligent Systems, 22*(2), 12–16.

Subrahmanian, V. S., Mannes, A., Sliva, A., Shakarian, J., & Dickerson, J. (2012). *Computational analysis of terrorist groups: Lashkar-e-Taiba.* New York, NY: Springer Publishing Company, Incorporated.

Tankel, S. (2011). *Storming the world stage: The story of Lashkar-e-Taiba.* London: C. Hurst & Co.

Tate, A., Levine, J., Jarvis, P., & Dalton, J. (2000, April). Using AI planning technology for Army small units operations. In *Proceedings of the Fifth International Conference on Artificial Intelligence Planning and Scheduling* (pp. 379-386). Edinburgh Univ. (UK): AAAI.

U.S. Army. (1997). *Field manual no. 101-5, staff organization and operations.* Washington, D.C.: Department of the Army.

Wagenhals, L. W., & Levis, A. H. (2007, April). Course of action analysis in a cultural landscape using influence nets. In *Proceedings of the IEEE Symposium on Computational Intelligence in Security and Defense Applications, 2007* (pp. 116–123). IEEE. Honolulu, HI.

Widener, M. J., Bar-Yam, Y., Gros, A., Metcalf, S. S., & Bar-Yam, Y. (2011). Modeling policy and agricultural decisions in Afghanistan. *GeoJournal 2013*, *78*(4), 591-599.

Yilmaz, L., Oren, T. I., & Ghasem-Aghaee, N. (2006, December). Simulation-based problem-solving environments for conflict studies. *Simulation & Gaming, 37*(4), 534–556.

Yost, K. A., & Washburn, A. R. (2000). The LP/POMDP marriage: Optimization with Imperfect Information. *Naval Research Logistics, 47*, 607–619.

Zhang, C., Cox, M. T., & Immaneni, T. (2002). *GTrans version 2.1 user manual and reference.* Wright State University, Department of Computer Science and Engineering. Dayton, OH.