Military Decision Aids—A Robust Decision-centered Approach

Patrick J. Talbot
TRW Systems

This article captures the results of research and development on aids for military decision making. The objective of this ongoing effort is to field a set of decision support tools for high-level decision making, capable of being applied across multiple mission domains. Our decision-centered approach is based on the Joint Operations Planning and Execution System (JOPES), which defines common decision processes for the U.S. Department of Defense. In our implementation, the decisions themselves are viewed as state transitions from current states to desired states according to a plan. Decision types span the full spectrum of military needs: situation diagnosis, pattern discovery, strategy determination, detailed planning, temporal reasoning, optimizing schedules, timeline enforcement, execution monitoring, and mission assessment.

Products to date include a robust development environment that has been applied to four mission areas. Most recently, a decision toolkit to integrate offensive and defensive missions was demonstrated. The knowledge acquired from our studies includes a concise understanding of decision types and conditions for which high-level military decision makers require automated support tools, as well as a deeper understanding of which reasoning and learning algorithms are useful in military decision making—and which are not.

Introduction

The motivation for this ongoing series of applied research and development projects is to demonstrate the potential utility of semiautomated, automated, and autonomous decision support tools applied to military missions. The effort is directed toward command and control simulation (C2Sim) prototypes, with eventual incorporation of proven decision support tools in operational environments. The foundation for joint military operations is the Joint Operations Planning and Execution System (JOPES) [1]. This framework provides broad guidance for the conduct of missions (in particular, when offensive action is required) in which joint forces are involved. We used JOPES to derive sets of military decisions. This decision-centered approach guided our development of displays and algorithms to support the decision maker.

The paradigm for defining a decision is action-oriented: a decision is a state transition from a current state to a desired state, according to a plan. The current state, desired state, and plan may or may not be known. Decision types span the full spectrum of military cognitive needs, such as situation diagnosis, pattern discovery, strategy determination, detailed planning, temporal reasoning, schedule optimization, timeline enforcement, trigger updates, execution monitoring, and mission assessment. The reasoning and learning algorithms are derived from the nature of the decisions required; for example, case-based reasoning for Course of Action (COA) and target determination, map-based planning for force allocation, an optimization policy for deciding when to update a plan, and Dempster–Shafer Belief

Networks for diagnosis and assessment. In addition to prototyping decision aids for operational systems, the research is also providing autonomous Simulated Commander's prototypes for C2Sims. Another direct benefit comprises the insights being provided to developers of operator displays.

Related work includes the cognitively motivated, Navy-sponsored Tactical DecisionMaking Under Stress [2] program. Multistrategy reasoning and learning is also evident in the Cognition-Oriented Emergent Behavior Architecture (COREBA) [3] testbed. The Defense Advanced Research Projects Agency has funded planning algorithms for the Synthetic Theater of War [4]; however, much of this effort used rule-based reasoning and was geared to low-level decision making. The Technology Assessments for Command Decision Modeling [5] are also germane. Much of the Computer-Generated Forces work that uses frameworks such as Suppressor [6], ModSaf/OneSaf [7], EADTB [8], and Soar [9] is complementary, but also addresses low-level decisions and is heavily rule-based. We produced a robust development process that was applied to four mission areas:

- Antisatellite mission algorithms
- A Command, Control, Communications, Computers, Intelligence, Surveillance, and Reconnaissance (C4ISR) demonstration
- A Simulated Commander for a C2Sim
- Offense/Defense Integration (ODI) of strategic deterrent forces (SDFs) and National Missile Defense (NMD) defensive forces with quantified force multiplier based on a new performance metric

Most recently, a decision toolkit to integrate Computer Network Defense and Computer Network Attack was addressed. Our contribution is a prototype testbed containing numerous decision-making tools and knowledge that we have acquired from our studies. This includes a concise understanding of decision types and conditions for which high-level military decision makers require automated support tools, and a deeper understanding of which reasoning and learning algorithms are useful in military decision making—and which are not.

Approach

Our top-down methodology (Figure 1) is decision-centered. Here, decisions required for ODI are derived from SDF and NMD mission areas. The methodology begins with an understanding of the mission domain, the concept of operations, and mission-related requirements. We leverage the JOPES to provide a common foundation for analyzing military processes. Decisions are defined as state transitions. Displays are based solely on organizing information and presenting it to a user in support of a decision. Reasoning and learning algorithms are formulated to populate the displays. Each of these steps is discussed in detail.

Integration of multiple missions requires a concept of operations that encompasses current practice, provides synergy among missions, and minimally intrudes on individual mission timelines. A spectrum of options (Figure 2) was identified to show that varying degrees of integration are possible. Tight integration of existing systems, shown as Single System, System of Systems, and Codependent Missions are desirable but not feasible because of differences in geographic location, mission requirements, prevailing culture, and integration costs.

Loosely coupled systems, such as Noninteracting Systems, Passive Observance, and Operator Dialog are undesirable because they lack synergy, may contend for scarce resources, and (in the

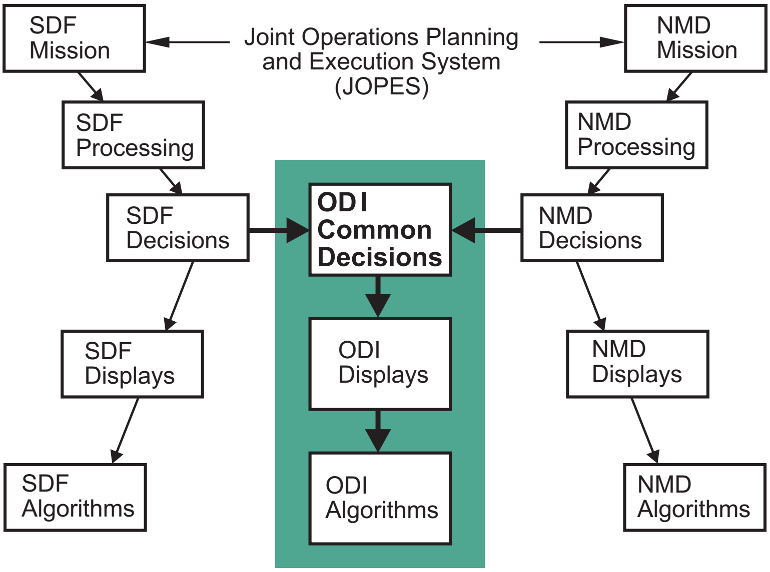

Figure 1. Decision-centered Methodology

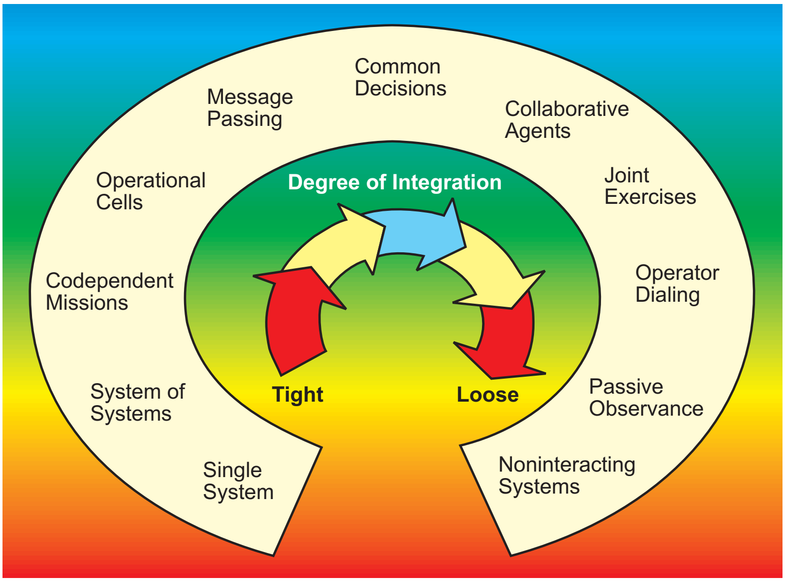

Figure 2. Spectrum of Options

extreme) may produce fratricide. Moving upward, Joint Exercises and Operational Cells are currently successfully employed and are very useful. An optimum integration includes Message Passing, Collaborative Agents, and Common Decisions, thus reinforcing the utility of a decision-centered process.

Common processes (Figure 3) are provided by the JOPES, thus ensuring reuse across military mission domains. Foreground processes defined by the JOPES are deliberate planning, mission planning, execution, and assessment. The feedback loop indicates the necessity for plan updates. The background tasks provide physics models, fog-of-war scenarios [10], and other utilities needed to provide context to the simulation environment.

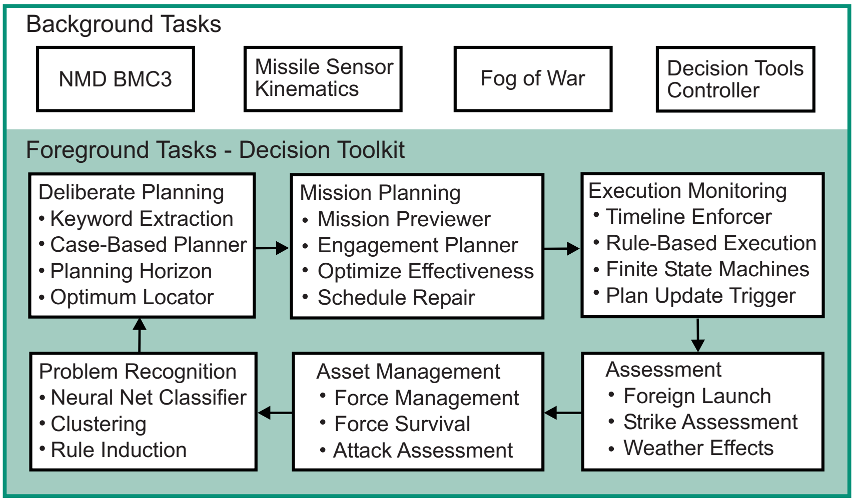

Figure 3. JOPES Processes

Decisions are defined as state transitions (Figure 4) from current states to desired states according to a plan. Current states, desired states, and a plan may or may not be known. Uncertainty in military decision making is a fact, not an afterthought. Eight combinations of known and unknown current states, desired states, and plans are possible. As a practical matter, problems with multiple unknowns, such as the initial intuitive problem (where nothing is known), are attacked piecewise. Four classes of military decision-making problems result: clarify current state, clarify desired state, determine alternatives, and execute. Full-spectrum decision tools were developed based on these four problems. These abstract problem types correspond to situation diagnosis, goal determination, detailed planning, and execution monitoring.

Displays have no purpose other than to support decisions. Extraneous content is an anathema to decision makers working in stressful environments and constrained by stringent deadlines. Because most military decisions are risky, uncertainty and conflict in data are shown explicitly. Our design goals were to use displays to support decisions in the following ways: information was organized hierarchically to hide detail; a windows-like look and feel was provided; the geographic, political, and historical contexts were available; animation provided a mission preview; and simple explanations were available.

Algorithms for reasoning and learning were formulated to help the decision maker indirectly. Rather than elevating fancy artificial intelligence techniques to a focus of the decision support process, the algorithms were structured to compute, organize, and enrich information for display.

Decision = State Transition

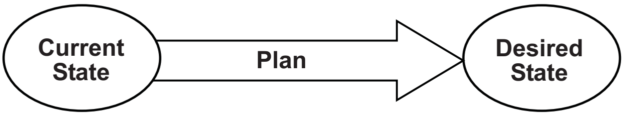

- Current State, Transition (Plan), and Desired States may or may not be known
- Paradigm has been shown to neatly sort decisions, algorithms, and displays

Representation	Decision Categories	Military Equivalent
●⇒○	Clarify Current State	Situation Diagnosis
○⇒●	Clarify Desired State	Goal Determination
○➡○	Define Alternatives	Detailed Planning
○⇒○	Execute Plan	Execution Monitoring

Legend: Known = ○ Unknown = ●

Figure 4. Decision as State Transitions

Testbed Demonstration System

Through a process of incremental refinement—over many projects and many years—we have evolved a rudimentary testbed to a reasonably robust, platform-independent, demonstration system (Figure 5). The domain is missile offense and defense. Work has begun on a testbed for information operations. The testbed features a Functional Model that computes the dynamics of simulation objects such as the Earth, satellites, missiles, and aircraft. Astrodynamic models calculate communications linkage, sensor look angles, and intercept missile guidance. Synchronized Parallel Environment for Emulation and Discrete-Event Simulation (SPEEDES) provides scalability by distributing processes over multiple processors to attain simulation speedup. A Situation Display provides a three-dimensional (3-D) animation, or mission preview. The Battle Planner schedules and tracks engagements—it is distinct from the Decision Aids module that provides reasoning algorithms and utilities, such as a fog-of-war overlay—to perturb the environment.

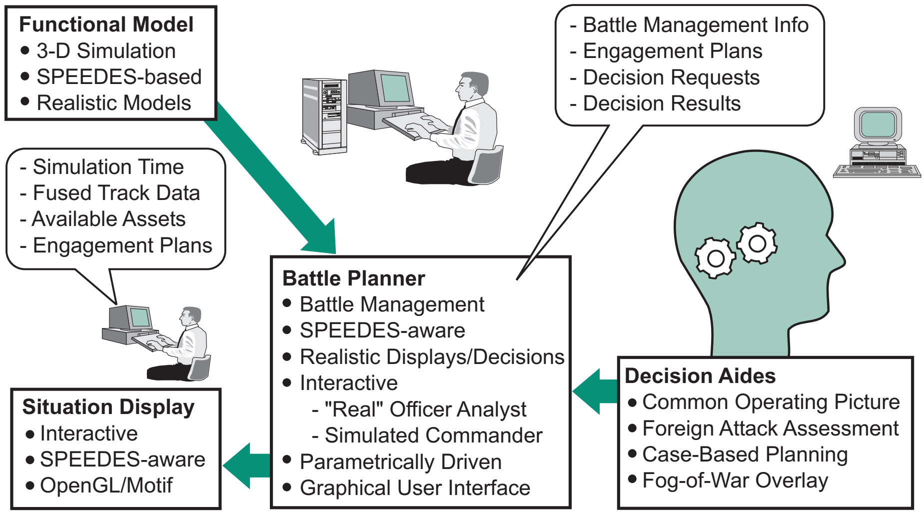

Figure 5. Demonstration Testbed

A needs analysis was completed that focused our efforts on the specific tasks where military decision makers were in need of help. The decision-making conditions, hard questions, and mission integration goals (Figure 6) formed a basis for ensuing efforts. Decision-making conditions can be very difficult, especially when cross-mission collaboration is required. The most challenging factors are uncertainty, workload, time pressure, and ill-posed tasks.

We further refined and synthesized our analysis as a reporter's checklist of basic questions. The "who, what, where, when, why, and how" (Table 1) were examined via sample problems that recur in all military mission areas. Human deficiencies were then identified. As might be expected, human cognitive abilities to answer the "who" question are excellent and need no decision support—people seem to know everything they need to about their roles and responsibilities and those of others around them. Conversely, humans do poorly with "when." As decision makers, we need reminders of deadlines, do not know when to abandon or repair a plan, and generally do not even feel that time flows linearly. Table 1 also maps these questions to problem types and the decision tools that we have developed to cope with these human deficiencies.

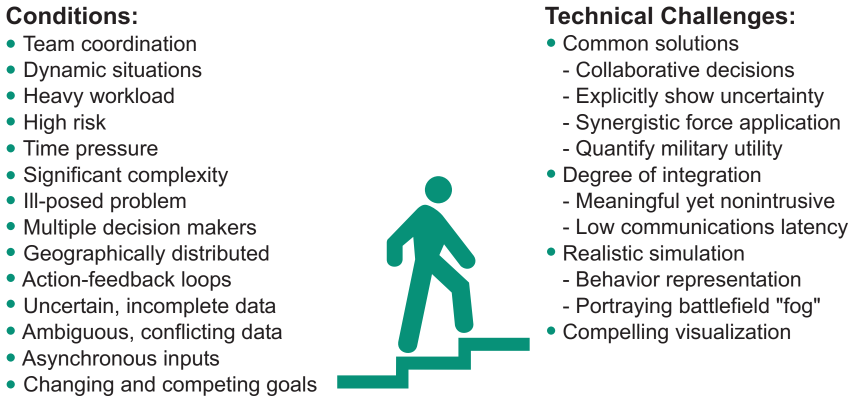

Figure 6. Decision-Making Conditions

Decision-Making Tools by Decision Types

Situation Diagnosis. This decision type is the foundation upon which other decisions are predicated. Proper diagnosis of situations is crucial because military decisions that adversely affect lives, property, and the environment are based on this assessment; hence, the decision is risky. The problem is characterized by uncertain, incomplete, ambiguous, and possibly conflicting data that arrive asynchronously. The difficulty is compounded because the time pressure is often intense. In these circumstances, people tend to refuse to decide, make hasty conclusions, ignore potentially important evidence, and have difficulty understanding data correlation.

Table 1. Human Decision-Making Strengths and Deficiencies

Questions	Sample Problems	Human Deficiencies	Decision Type	Decision Tools
Who	Who's in charge, roles, responsibility, interfaces	None, humans are extemely proficient	Situation diagnosis, execution monitoring	None required
What	Assess threat, response goal, constraints, strategy	Handling uncertainty, data overload, limited memory	Situation diagnosis, goal determination	Data fusion, case-based planner, situation map
When	Event timing, task duration, time constraints, plan update	Few cognitive skills, time does not flow evenly, difficult to reprioritize under time pressure	Situation diagnosis, goal determination, execution monitoring, detailed planning	Timeline enforcer, planning horizon, mission previewer, assessment deadline, replan trigger, optimizer
Where	Weapon placement, target determination, big picture	Few, highly proficient, memory overload	Situation diagnosis, detailed planning	Common operating picture, map-based planning
Why	Enemy intent/ strategy, interpretation of broad goals	Handling uncertainty and conflict in evidence	Situation diagnosis, goal determination	Data fusion, case-based planner
How	Engagement plan, schedule, logistics, process	Training provides proficiency, optimization is difficult	Detailed planning	Battle management, schedule optimizer
How Much	How good is the plan, how far ahead to plan, how many resources to commit	Qualifying probability of success, handling uncertainty, determining robustness	Detailed planning, execution monitoring	Figures of merit, schedule repair, schedule optimizer

Our solution to situation diagnosis is a Dempster-Shafer Belief Network [11]. Specific applications that have been demonstrated are Anti-Satellite strike assessment, NMD strike assessment, space weather effects, and foreign launch assessment (Figure 7). The advantage of this formulation is that it explicitly shows belief and plausibility (degree of conflict) in hypotheses, shows how these increase or decrease over time, and provides an explanation. Also shown explicitly in the plot are decision deadlines and confidence thresholds.

Goal Determination. The problem is to decide on a goal state; that is, *what* we want to achieve. The characteristics are:

- A decision maker is given broad objectives that result in an ill-posed problem
- There is an exponential increase in possible solutions, having many dimensions or criteria
- It is difficult to quantitatively rate the worth of a strategy

When faced with these circumstances, people often attempt to recall a course of action that worked under similar circumstances; caucus to arrive at a consensus (group-think); or because of cognitive overload, identify a dominant criterion and assemble a strategy around it.

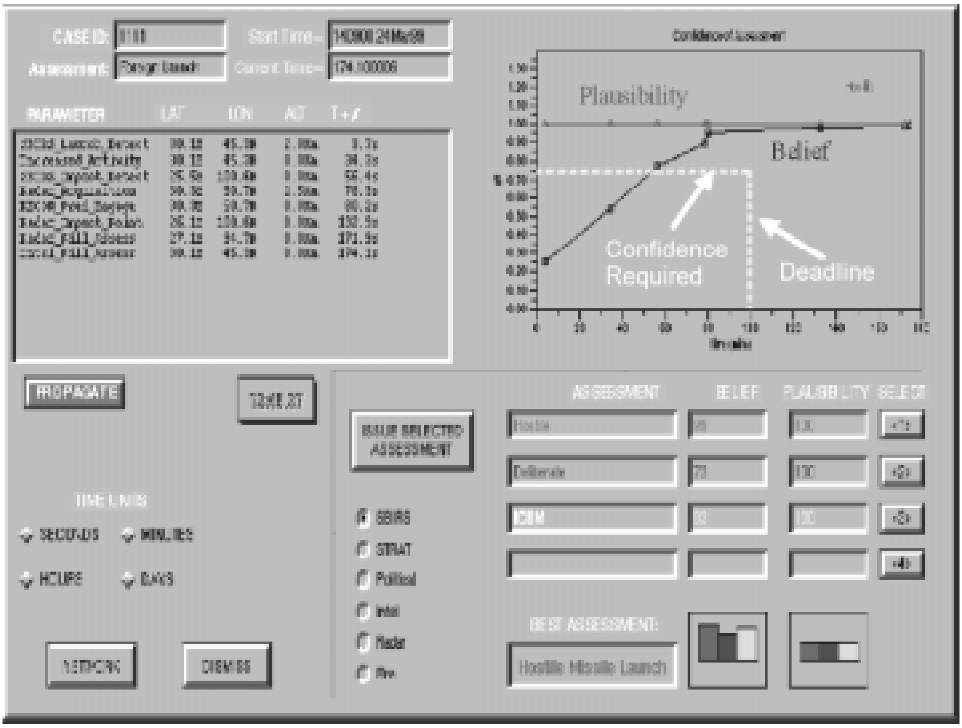

BEST ASSESSMENT:

Hostile Missile Launch

Figure 7. Dempster-Shafer Belief Network for Foreign Launch Assessment

Our solution to goal determination is a case-based planner (Figure 8) that allows preplanned options to be retrieved, scored for similarity to weighted selection criteria, and presented for the decision maker to edit, since it is easier to edit than create. Likening case-based reasoning to an engineering trade study (where options are compared to one another based on weighted selection criteria) provided sufficient insight into the algorithm to quickly implement an application. This process helps an operator select COAs for inclusion in a Commander's estimate. We also implemented an Optimum Locator which combined case-based planning, rule-based reasoning, and astrodynamic models to decide where best to put offensive resources based on kinematic and logistic constraints. Target determination [12], a related problem, is in work. Advantages to this technology are that it mimics human decision making, produces a finite set of cases, computes a score, and provides an explanation. It is also a good scheme for schedule repair. Figure 8a shows a hierarchically structured input screen for selecting and weighting COA criteria. The highest scoring result—computed as a weighted sum of the words in the incoming Joint Chiefs of Staff Warning Order that match words in stored cases—is shown pictorially in Figure 8b. A detailed, color-coded comparison of selection criteria and cases is shown in Figure 8c.

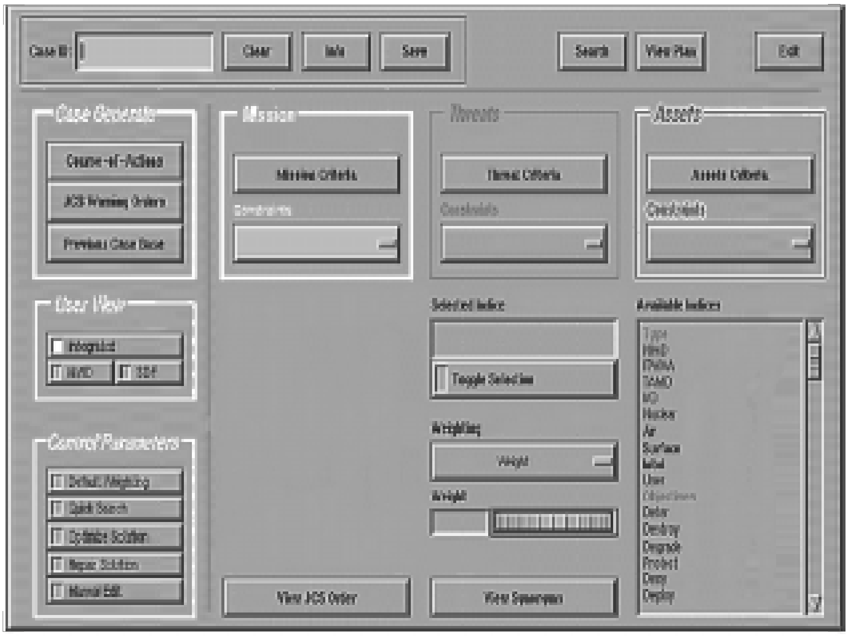

a. Input Screen for Selecting COA Criteria

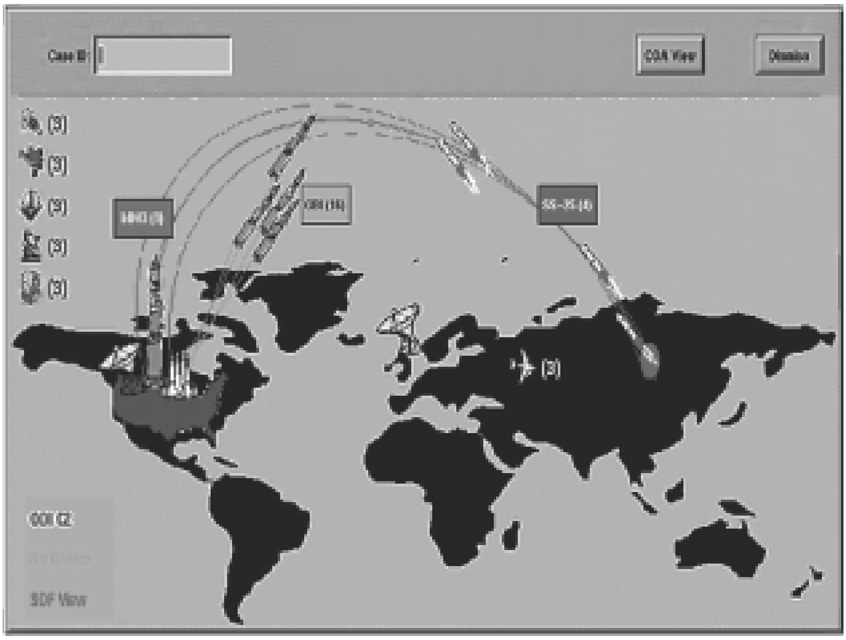

b. A Map View for Visualizing Combined Planning Results

Figure 8. Case-Based Planner for Course of Action Determination

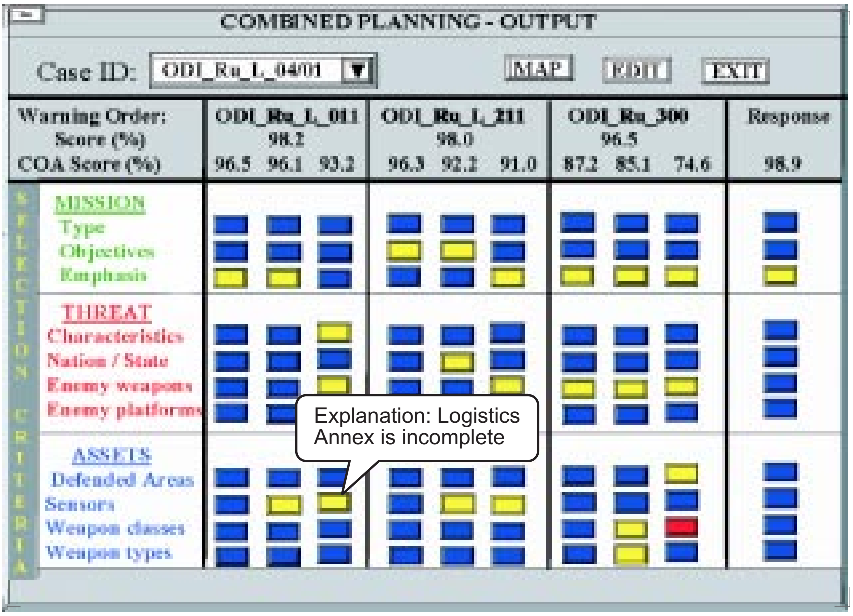

Case ID: ODI_Ru_L_04/01 ▼

Warning Order:	ODI_Ru_L_011			ODI_Ru_L_211			ODI_Ru_300			Response
Score (%)	98.2			98.0			96.5			
COA Score (%)	96.5	96.1	93.2	96.3	92.2	91.0	87.2	85.1	74.6	98.9
MISSION										
Type	🟦	🟦	🟦	🟦	🟦	🟦	🟦	🟦	🟦	🟦
Objectives	🟦	🟦	🟦	🟨	🟨	🟦	🟦	🟦	🟦	🟦
Emphasis	🟨	🟨	🟦	🟦	🟦	🟨	🟨	🟨	🟨	🟨
THREAT										
Characteristics	🟦	🟦	🟨	🟦	🟦	🟦	🟦	🟦	🟦	🟦
Nation / State	🟦	🟦	🟦	🟦	🟨	🟦	🟦	🟦	🟦	🟦
Enemy weapons	🟦	🟦	🟨	🟦	🟦	🟨	🟨	🟨	🟨	🟦
Enemy platforms	🟦	🟦	[illegible]	[illegible]	[illegible]	[illegible]	🟦	🟦	🟦	🟦
ASSETS										
Defended Areas	🟦	🟦	🟦	🟦	🟦	🟦	🟦	🟦	🟨	🟦
Sensors	🟦	🟨	🟨	🟦	🟨	🟨	🟦	🟦	🟦	🟦
Weapon classes	🟦	🟦	🟦	🟦	🟦	🟦	🟦	🟨	🟥	🟦
Weapon types	🟦	🟦	🟦	🟦	🟦	🟦	🟦	🟨	🟦	🟦

Explanation: Logistics Annex is incomplete

c. Combined Planning Detailed Output

Figure 8. Case-Based Planner for Course of Action Determination (Concluded)

Detailed Planning. The problem was to elaborate a selected COA in sufficient detail to produce a schedule of activity. This was comparable to the "Job Shop Scheduling Problem," but was complicated by the dynamic nature of warfare and the inherent uncertainty associated with the concurrent execution of the schedule. Case-based planning was again employed (for the reasons cited previously) to produce both an engagement plan and schedule repair. It was augmented by a genetic algorithm to optimize the schedule. A Mission Previewer provides a 3-D animated view of the sequence of actions being scheduled. The latter map-based planning has had a long history of success in military decision making. To answer the "how much" question and deal with the underlying uncertainty, we prototyped a Planning Horizon algorithm that used the two-dimensional Monte Carlo sampling theorem to determine how far ahead we can feasibly schedule an activity before uncertainties accumulate sufficiently to undermine the effort.

Execution Monitoring. Given that we had clarified the current and desired states and had a plan, the problem was then to provide decision-making tools to execute the mission and monitor progress. We used rule-based reasoning for mission execution. A small number of rules (about 90) sufficed for the NMD mission. Given the precise nature of military doctrine applied to a limited number of execution decisions (eight for NMD), this was reasonably straightforward. Multiple challenges arose for the decision maker in other areas: translating low-level status inputs to mission availability, and many decisions related to the "when" question. We demonstrated two tools for assessing mission availability. Our low-technology approach was to use a simple matrix multiplication scheme to compute an estimate. We also used a knowledge discovery tool [13] to find the underlying structure (Figure 9) in a STRATCOM Force Management database using a rule tree induction algorithm (a rule induction tree).

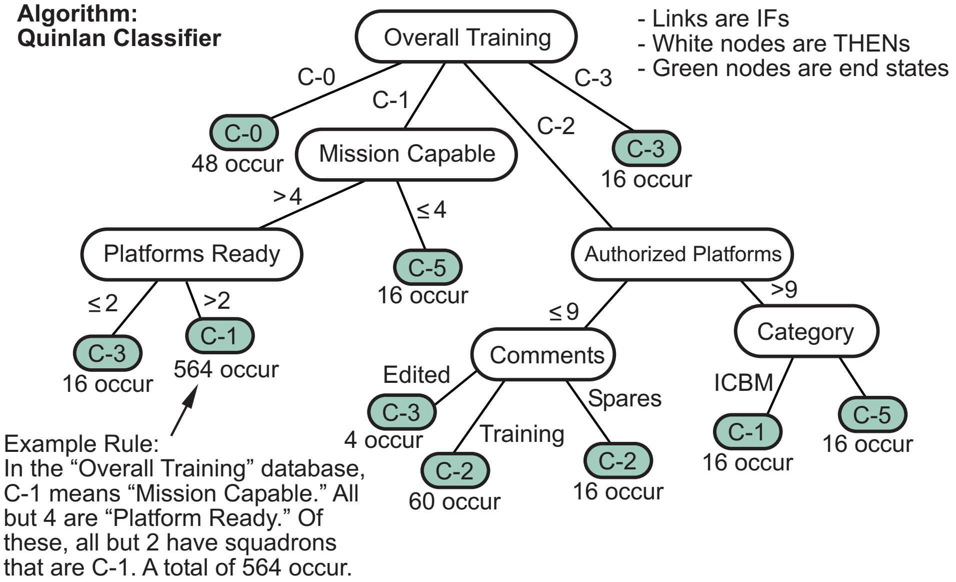

Figure 9. Rule Induction Tree for READI Force Management Data

The "when" decisions required temporal reasoning, and were more varied and challenging. To date, we have produced a timeline enforcer that provides a view of deadlines and alerts the operator as they approach. A particularly difficult problem for military decision makers is to decide when to repair or to replace a plan, given that it is going (or predicted to go) bad. Faced with these circumstances, people often (1) abandon the plan and continue prosecuting the mission in an ad hoc fashion; (2) continue with the plan for too long; or (3) cannot decide whether to repair or replace the plan. They generally do not seem to have a cognitive strategy for coping with the problem.

Our solution (Figure 10) to help a decision maker decide when to repair or replace a plan is an optimal policy algorithm [14] that computes the best time to repair a detailed plan. The decision point is based on a long-term strategy that minimizes the ratio of planning time to execution time (ordinate). It is a function of mission time (abscissa), probability of success of the current plan, the time required to repair or replace the current plan, and the anticipated probability of success of the new plan.

Results

Based on the lessons we learned by developing and demonstrating various decision tools, and the peculiarities of military decision making, we provide an assessment of the technology (Figure 11) and what works and what does not. The most effective algorithmic processing schemes reflect our current focus.

For planning, case-based reasoning is exceptionally powerful, flexible, and acceptable to decision makers. That is the way they often devise strategies, reuse previous plans, and identify plan repairs. Simulation is an excellent way to test a plan in a virtual space without actually executing it. Constraint satisfaction is built in to case-based reasoning selection criteria with *must-have*, and *not* logic. Genetic algorithms are good auxiliary techniques for plan optimization, but are not a primary means of planning; and operations research requires a well-posed problem—which is usually not the case.

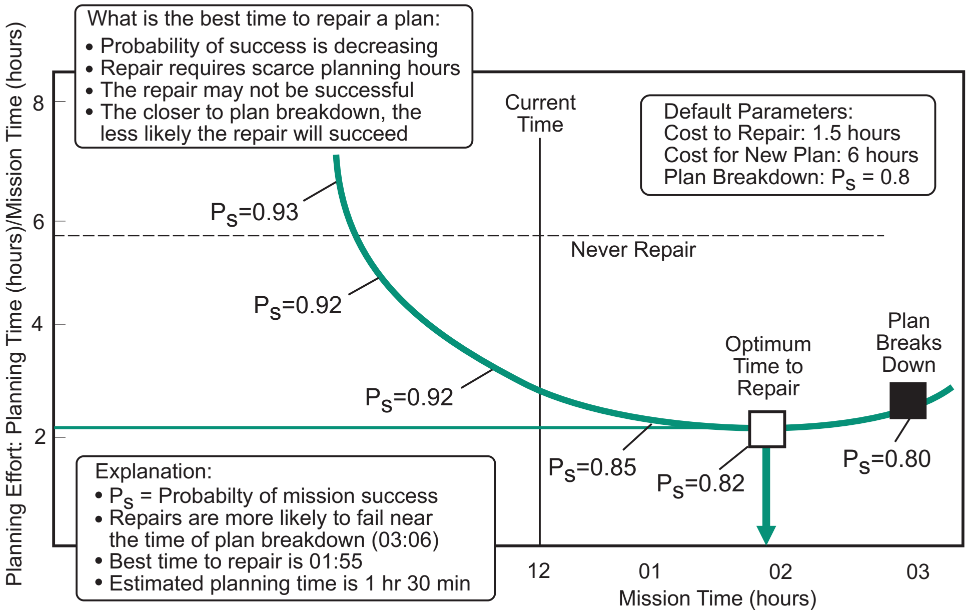

Figure 10. Replan Trigger

	Planning	Visualization	Data Fusion	Learning	Reasoning
Effective ↑	Case Based	Map Based	Belief Networks	Rule Based	Heuristic **FOCUS**
	Simulation	Animation	Heuristic	Case Based	Optimal Policy
UTILITY	Constraint Satisfaction	HTML	Templating	Fuzzy Logic	
	Genetic Algorithms	Intelligent Multimedia Interfaces	Classical Statistics	Genetic Algorithms	Case Based
↓ Ineffective	Operations Research	Learning Agents VR	Rule Induction	Neural Nets	Fuzzy Logic

Figure 11. Techniques Assessment

For visualization, map-based planning with animation is extremely well-suited to military operations. Web-based content delivery (HTML) is still plagued with latency. Intelligent multimedia interfaces [15] that allow a response to be fashioned on the fly may be viable soon. Finally, virtual reality has promise, but higher-level military decision makers do not like to wear goggles. The advent of autostereoscopic displays will make immersive human-machine interaction more palatable.

Data fusion is an important element in situation diagnosis, and belief networks (we used Dempster-Shafer) have worked well for us. For military applications, rules are readily derived from doctrine, rules of engagement (ROEs), and common sense, and help reduce the sizes of the networks. Templating, classical statistics, and rule induction all require large amounts of evidence (which typically are not available at decision time) and are too brittle.

Learning rules and cases are straightforward and powerful. Fuzzy logic (we used Fuzzy CLIPS) was not necessary because of the precise nature of military doctrine, ROEs, and commands. Genetic algorithms are overkill, except for schedule optimization, and neural nets do not learn precise patterns nor do they typically provide adequate explanations.

Reasoning based on heuristics (rules of thumb) enabled us to keep search spaces fairly small. The rule-induction algorithm helped us to understand the underlying data structure in a force management database. We also used rule induction to check handcrafted rule sets for consistency and completeness: the algorithm learned a rule tree whose structure immediately showed problems. Optimal policy was useful for temporal reasoning. Again, case-based reasoning was only good for planning, and fuzzy logic was available but not needed.

Future Directions

To date, we have not had explicit need for agent technology because we operate in a homogeneous, well-understood, and controlled environment. However, intelligent software agent technology appears promising for collaborative planning—specifically, for coherently combining and deconflicting offensive and defensive partial plans.

The artificial life community has made significant strides in demonstrating emergent behavior in open systems [16] by evolving populations of simulated entities who are imbued with attributes and interact with the synthetic environment based on rules. This paradigm shows promise for understanding military deterrence and the interactions that trigger escalation in strategic offense/defense integration; for example, what is the effect on Arms Control if the United States fields a National Missile Defense system?

Finally, we plan to apply Minsky's "Society-of-Mind" concept to military decision making. The premise is that multistrategy reasoning techniques provide opinions that can be fused to provide robustness in the presence of fog of war.

Summary

Several unique and innovative aspects of decision support algorithms have been discussed. The foundation was the JOPES process, which fostered broad applicability to many military domains. A decision was succinctly defined as a state transition from a current state to a desired state according to a plan. The inherent uncertainty in decision making was explicitly considered in the formulation of decision types, not "layered-on" as an afterthought. Finally, our experience with a wide variety of algorithms from the fields of engineering, artificial intelligence, and statistics was discussed, and a subjective ranking of techniques based on effectiveness in military decision support was provided.

References

[1] *Joint Staff Officer's Guide AFSC,* Pub. 1, 1997, *http://www.afsc.edu/pub1/afsc0000.htm.*

[2] J.G. Morrison, R.T. Kelly, R.A. Moore, and S.G. Hutchins, "Tactical Decision Making Under Stress (TADMUS)—Decision Support System," *1997 IRIS National Symposium on Sensor and Data Fusion*, April 14–17, MIT Lincoln Laboratory, Lexington, Massachusetts, 1997.

[3] S. Kwak, "Cognition-Oriented Emergent Behavior Architecture (COREBA)," *Proc. SPIE—AeroSense Conf.*, Vol. 4026, March 2000.

[4] *http://stow98.spawar.navy.mil*, as of March 2001.

[5] *http://www-leav.army.mil/nsc/warsim/reason/assess/section2/section.htm/* as of March 2001.

[6] *www.afsaa.hq.af.mil/SAA/MODELS/suppressor.shtml* as of March 2001.

[7] *http://www-leav.army.mil/nsc/stow/saf/onesaf/* as of March 2001.

[8] *http://www.smdc.army.mil/EADTBTEX.html.*

[9] C.B. Congdon and J.E. Laird, *The Soar User's Manual: Version 7,* Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, Michigan, 1995.

[10] P. Talbot, "Simulating Fog-of-War in Military Decision Making," presented at International Conference on Industrial & Engineering Applications of Artificial Intelligence & Expert Systems (IEA/AIE 2000), New Orleans, Louisiana, July 2000.

[11] D. Zarley et al., "Evidential Reasoning Using DELIEF," *Readings in Uncertain Reasoning*, Morgan Kauffman, San Mateo, California, 1990.

[12] D.E. Moriarity, "Determining Effective Military Decisive Points through Knowledge-Rich Case-Based Reasoning," *Lecture Notes in Artificial Intelligence 1821*, Springer, 2000.

[13] I.H.Witten and E. Frank, *Data Mining*, Morgan Kaufmann, San Francisco, California, 2000.

[14] S.M. Ross, *Applied Probability Models with Optimization Applications*, Dover, New York, 1970.

[15] M.T. Maybury (ed.), *Intelligent Multimedia Interfaces*, AAAI Press/The MIT Press, Menlo Park, California, 1993.

[16] D.A. Fisher and H.F. Lipson, "Emergent Algorithms—A New Method for Enhancing Survivability in Unbounded Systems," *Proc. 31st Hawaii Intl. Conf. Syst. Sci.*, Maui, Hawaii, January 6–9, 1998.

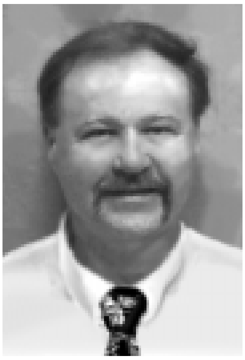

Patrick J. Talbot has extensive experience with TRW, primarily in Independent Research and Development. He is currently the principal investigator on the Offense/Defense Integration IR&D project. He has performed on Minuteman, spacecraft projects, battle management tasks, antisatellite systems, and the development of decision algorithms. In addition, he has managed a Delivery Order titled "Technology Insertion Studies and Analysis" at the Joint National Test Facility in Colorado Springs, Colorado, where his research included parallel computing, simulated human behavior, and strategic planning. He holds a BS in applied mathematics, a BS in physics, and an MS in physics, all from Pennsylvania State University.

e-mail: Pat.Talbot@trw.com