

experts) to build adversarial decision models, given the non-cooperative nature of the adversary?

4. How do we use these models to support friendly decision makers and to increase their predictive battlespace awareness?
5. How do we test and evaluate our models to determine how well they perform in a broad range of situations?

We are addressing the issues listed above through an Information Institute Research Project (IIRP) funded by the Air Force Research Laboratory's (AFRL) Information Directorate and performed in collaboration with AFRL's Human Effectiveness Directorate. Specifically, we are investigating those salient human factors characteristics that must be modeled (issue #1); the creation of efficient and effective computational models (issue #2) that, given observations of an adversary's actions and reactions (issue #3), generates hypotheses about the adversary's intent and suggests appropriate responses (issue #4). The efficacy of our models within cooperative domains has already been proven (Bell, Franke, and Mendenhall, 2000; Franke, et al. 2000). The value-added of our models to Department of Defense (DoD) personnel performing adversary course of action prediction will be iteratively evaluated during the IIRP (issue #5).

Adversarial decision modeling technology yields ideas about what the adversary is trying to accomplish, as well as explanations about why the adversary is trying to accomplish those particular objectives. Deriving hypotheses about future actions of an adversary requires *information* about the adversary's current actions and *inferences* about the adversary's motivations. The informational requirements can be approached by bringing together knowledge elicitation, data collection and data fusion capabilities. The inferential requirements can be approached by creating models to both generate *descriptive* probabilities (to what extent does motivation X account for the set of observations Y?) and *predictive* probabilities (how likely is future action Z given motivation X?). The results of this system can be communicated to situation assessment tools to further refine the overall operational picture within a particular tactical setting. Further discussions of the adversary models and information collection and fusion issues can be found in Bell, Santos Jr. and Brown (2002).

Background

Intelligence Preparation of the Battlespace and COA Development

A key process used by the United States military to predict adversary courses of action is the Intelligence Preparation of the Battlespace (IPB) process. The goal of this doctrinally driven, four-step process is to "...reduce uncertainties concerning the enemy, environment, and terrain for all types of operations." (Joint Publication 2-01.3 2000). This is achieved by determining the adversary's likely COA, de-

scribing the environment friendly forces are operating within and the effects of this environment on these forces ability to achieve their goals. The four-step process is briefly provided below:

- 1. Define the battlespace:** assess the crisis situation; review commander's guidance / objectives; identify limits of operational area, area of interest, significant battlespace characteristics; evaluate existing databases and identify information gaps; obtain products / information required to conduct remainder of IPB
- 2. Describe battlespace effects:** describe how characteristics of the surface, aerospace, information, human and weather dimensions affect operations employment; identify information gaps
- 3. Evaluate the adversary:** map relevant adversary processes and identify friendly and adversary centers of gravity, capabilities, limitations and vulnerabilities; perform critical nodes analysis to identify high value targets; identify information gaps
- 4. Determine adversary COAs:** identify, evaluate and prioritize the adversary's likely objectives and desired end state and the full set of COAs available to the adversary while identifying initial information collection requirements

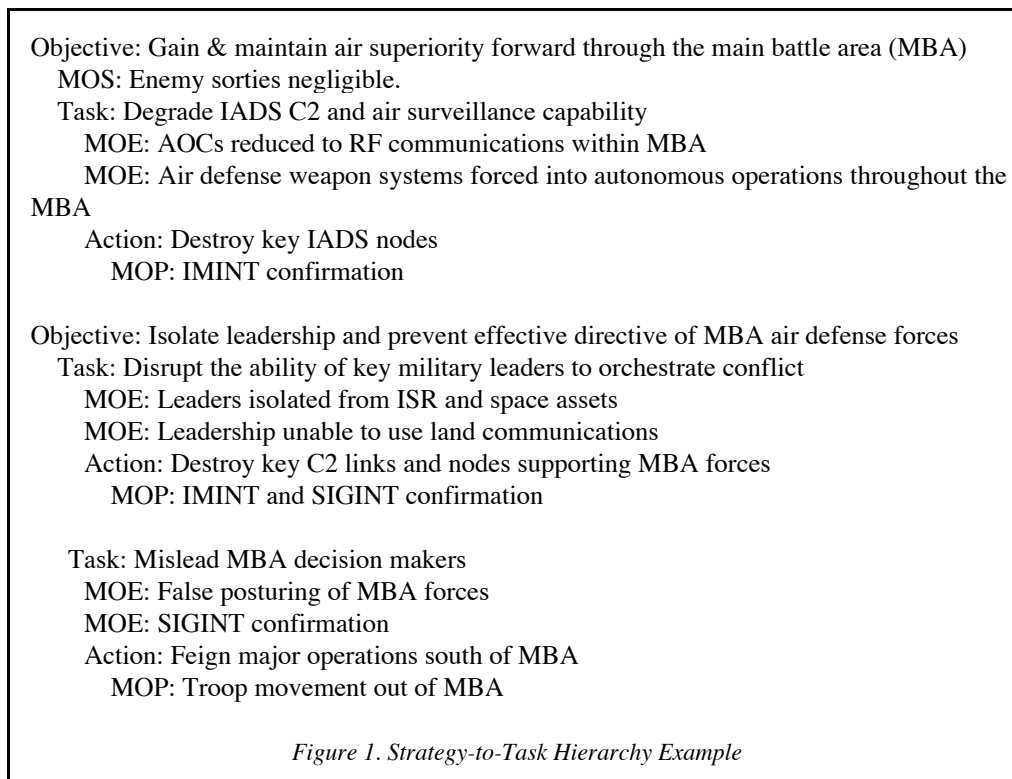
The IPB process, as part of operational environment research, is a major input into (both friendly and adversary) objective determination for deliberate and crisis action planning, and to a lesser extent, force execution due to the time and information demands of the IPB process.

COA development¹ begins with reviewing Combatant Commander guidance, intent and objectives and then constructing a strategy-to-task framework. This framework allows planners to determine supporting objectives (e.g., service component objectives that support the commander's objectives), tasks, actions, and targets based on the IPB information they have previously generated on a particular adversary.

For each level in the strategy-to-task decomposition hierarchy, success indicators (sometimes called battle damage indicators or measures of success (MOS)²) are assigned as observables with quantitative and / or qualitative metrics. During this COA development-planning phase, the IPB information is refined and additional information is collected.

¹ There are two COA development activities being performed simultaneously—one to determine how our time-phased actions will meet a commander's objectives and one to determine what actions an adversary might take to achieve his objectives. Operational planners typically do the former while intelligence analysts do the latter. To avoid confusion, we will always fully designate the adversary COA development process; otherwise, we will mean the friendly COA development process. To a large extent, the products generated are the same.

² Measures of success attach to a desired end state (that is, what the situation should look like once the operation is over). Since the objective is typically accomplished over time, the end state is decomposed into a series of phased events. An event should have a MOS. Sets of conditions determine whether the event has occurred. Measures of effectiveness and performance (MOEs and MOPs) attach to conditions.



An example of a fictitious strategy-to-task hierarchical decomposition is provided in Figure 1 below. Numerous tools support this strategy-to-task decomposition including AFRL's Effects-Based Operations Advanced Technology Demonstration Strategy Development Tools, Joint Information Operations Center's Information Operations Navigator, and Electronic Systems Center's Information Warfare Planning Capability.

It should be noted, however, that few tools exist that explicitly capture the adversary's COAs. A notable exception is the Target Prioritization Tool for Links and Nodes (TPT-LN). The use of a strategy-to-task hierarchy to represent adversary COAs has merit for military planners. The advantages include re-using a well-known process and representation (i.e., a hierarchical decomposition) and the ability to explicitly compare and contrast friendly force COAs with competing adversary COAs.

Effects-Based Operations and Adversary Intent Inferencing

Target-based and objectives-based (the well-known "strategy-to-task" approach described above) approaches to planning do not explicitly address the adversary's decision-making processes. A new approach to planning that explicitly addresses the adversary must be developed. The basis for such an approach is emerging from USAF-sponsored research. This approach, termed effects-based

operations (EBO), is the best candidate to serve as the basis of the operations model we require (McCrabb 2000). Basically stated, EBO is "an approach that...explicitly seeks to understand, trace, and anticipate direct and indirect effects of a specific action...on an adversary's course of action." (Fayette 2001) EBO is framed with respect to outcomes produced (and / or predicted to be produced) in the battlespace. *EBO inherently addresses an adversary as a system.* The notion of "effect" is predicated upon the presumption that there is an object of reference (specifically one systemically organized), namely the adversary, whose state(s) can be identified and influenced through prospective courses of action. EBO planning is predicated on a coherent model of the state(s) and dynamics of the adversary system(s). At the center of the EBO concept is the idea that effective friendly COA planning can and should be framed with respect to effects to be induced in an adversary system.

The key to effects-based operations revolves around determining how an adversary *should / can / could* react to system perturbations resulting from actions on the battlefield from our own forces (McCrabb 2000). One of the greatest technological challenges for the EBO approach is that of adversarial decision modeling. While EBO's overall goal is to model the enemy in its entirety (stated as "enemy-as-a-system" in the EBO CONOPs and including the physical, data, cognitive, and social aspects of the battlespace centers of gravity and the dependency linkages be-

tween them), we believe that a necessary starting point is to model an *adversary commander's intent*. Intent¹ inference involves deducing an individual's goals based on observations of that individual's actions (Geddes 1986). In automated intent inference, this process is typically implemented through one or more behavioral models that have been constructed and optimized for the individual's behavior patterns. In an automated intent inference system, data representing observations of an individual, the individual's actions, or the individual's environment (collectively called *observables*) are collected and delivered to the model(s), which match the observables against patterns of behavior and derive inferred intent from those patterns. These inferences can then be passed to an application for generation of advice, definition of future information requirements, proactive aiding, or a host of other benefits. Furthermore, the success of adversary intent inferencing addresses a key technological barrier of EBO—that of the human element's impact in EBO. Once adversary intent is suitably modeled and captured, we can then compose these individual adversary commander's intent models into larger collectives using our work in team intent modeling (Franke, et al. 2000) to address the general problem of the “enemy-as-a-system”.

A Knowledge Acquisition Approach

In this section, we detail our approach for performing the initial and subsequent knowledge acquisition to build adversarial decision models. Our approach is pragmatic in its use of pre-existing processes, tools, and data already in wide use by the DoD planning community. The construction of models to support adversary intent inferencing will be driven by a number of different sources. The single most influential source of adversary intent modeling information must be the human subject matter experts (SMEs) who are most familiar with particular adversaries. Doctrinal knowledge can provide a foundation on which to build, but does not offer a complete solution. These SMEs can provide our models with the intuitive reasoning that cold facts and rigid doctrine cannot. In addition, military planners must also address historical case studies of the adversary and up-to-date information of the political environment of the region(s) in question. By providing intuitive means for an SME to specify possible adversary objectives, the relationships between these objectives and the tasks and actions that constitute an adversary's courses of action, an adversary intent inferencing system can provide assistance with collecting and organizing command and control information (Hofmann et al., 2000; McGrath, Chacon, and Whitebread, 2000), gathering and monitoring in-

¹ What exactly constitutes intent has long been debated in the cognitive and psychological sciences (as well as artificial intelligence). For purposes of this discussion we stand on the following military-oriented definition: Intent is composed of a commander's desired end-state/goal, the purpose/reason for pursuing that end-state, a methods/means to achieve the end-state, and a level of commitment to achieving that end-state (based on acceptable risk of the pursuit and probability of success).

formation in intelligence databases (Whitebread and Jameson 1995) and performing multiple mission planning and execution activities (Saba and Santos Jr. 2000). One advantage to our pragmatic approach is that it allows for an incremental, phased approach to adversary course of action prediction. We fully realize that a model is only as good as the data that supports that model. As any particular situations “flares up” and military planners start the IPB process for an adversary, developing related intelligence for the area of interest and therefore learning as they go along, given there is little to no existing information on a given adversary.

Mapping the User's Domain to Adversary Intent Inferencing

While observables in the user intent domain stem from data collected from use of computer systems by humans, observables in the adversary intent domain take the form of tactical information derived from intelligence databases, observations of the tactical environment, and input from human experts interacting with the adversary intent models. In place of window events, keystrokes, and mouse movements common in the user intent domain, our system in the adversary intent domain uses information about adversary location, movements, and activities to drive its intent inference processes. In place of computer state, analyses of information queries, and the content of user dialogue with team members, our system bases inferences on facts about the local terrain, regional weather, and the salient political climate.

The modeling process begins by analyzing military planners who are performing the Intelligence Preparation of the Battlespace (IPB) process. These planners define the range of objectives (i.e., end-states or goals) that an adversary might attempt to carry out and the available actions (i.e., means / methods) that the adversary has for carrying out those objectives. These objectives and actions represent the space of possibilities that the intent inference system must explore in examining adversary behavior. The modeler must also identify the observables associated with each individual action and indicate the method by which each observable can be ascertained. This will guide the integration of the intent inference mechanism into an operational context.

Next, the modeler must choose appropriate system architectures to capture the relationships between objectives, actions, and the observables. Given the simple fact that intent inference is an inherently uncertain process, this system architecture must both be able to deal with uncertain (and incomplete) information, as well as adapt over time as the result of new (possibly previously unknown or unanticipated) information. Finally, the modeler must decide upon the expected use of the system and implement a surrounding framework to make use of the models' inferences. Should the system perform descriptive, predictive, or diagnostic functions? How will the intent inference system provide timely, beneficial assistance to a decision maker?

We address several aspects of the system architecture design, including knowledge acquisition using the IPB process, the adversary intent inferencing model and the collection and production of observables, in the following sub-sections.

Knowledge Acquisition Using The IPB Process. The Intelligence Preparation of the Battlespace (IPB) process, first presented in the Background section above, is used by military intelligence analysts to arrive at an estimation of the adversary's possible courses of action (COAs). The IPB process provides an excellent basis for our knowledge acquisition. The output of the last step of the IPB process is very similar to the output from our adversary course of action prediction system, namely an estimation of the intent of the adversary and how the adversary is likely to act into the future. By reviewing the current IPB process, conducted primarily by human analysts with only limited automation support, we believe that it will be possible to determine much relevant information for creating a functional adversarial decision modeling (ADM) system.

First, it is possible to enumerate the various data sources that military intelligence analysts typically access in order to perform the IPB process. The structured data sources have schema that can be reviewed to create lists of data fields, and the unstructured data sources can be noted as possible candidates for automated evidence extraction and link discovery. In addition, analysts are likely to have some common tactical picture display in front of them, driven by data fusion processes. Such data fusion outputs would also be provided to an ADM system. Analysis of the IPB process allows us to enumerate the inputs most likely needed to support automated ADM.

Second, IPB analysts combine the evidence provided to them in certain ways in order to arrive at a hypothesis of the adversary's intent. Operators begin this process by analyzing the adversary from a number of different perspectives. Perspectives include political, cultural (research being performed into so-called "cultural lens" lends credence to the idea that culture should be considered when modeling adversaries), personality, emotional, economic, technological, will-to-win, risk perception, fatigue and morale. For example, if the adversary forces have not received supply within a certain period of time, and have been almost continuously under attack, it is likely that their morale will be low. Likewise if enemy forces are only weakly allied with their leaders, they may not share the same will-to-win. IPB analysts perform a series of such reasoning steps within a prescribed set of perspectives. By enumerating these perspectives, we believe that it will be possible to ensure that an ADM system, employing and reasoning about many of these same perspectives, will perform analysis of adversary intent in a comprehensive and exhaustive fashion. The Air Force Research Laboratories has been researching a number of descriptive decision models for adversarial decision-making (See, and Kuperman 1997; Llinas, Drury, Bialas, and Chen 1998; Llinas, Drury, Jian, Bisantz, and Younho Seong 1999).

Third, once IPB analysts have collected data, combined it into evidence under one of the several analysis perspectives, they infer adversary intent and determine plausible adversary courses of action. In order to do this, analysts, through training and experience, develop rules to map observables of adversary actions and general background data (such as fact books on the adversary's country and culture) to categories of intent. At the broadest level such categories of actions may be advance, attack, retreat or defend. With further refinement, it is possible to say, for example, that destroying a particular bridge is the intent of the adversary under the broader category of attack. These rules, mapping inputs to predictions of adversary intent are the most valuable aspect of the IPB process to capture and distill for use in ADM. We recognize that although some of these rules are made explicit in the IPB process, many are refined over the course of time by human operators and will thus be harder to capture and incorporate in the ADM system.

Adversary Intent Inferencing Model. The components of our adversary intent inferencing model, and the interactions between these components, are shown in Figure 2 below. The three core components that comprise our architecture and functions are as follows:

1. Goals: Prioritized short- and long-term goals list, representing adversary intents, objectives or foci
2. Rationale: A probabilistic network, representing the influences of the adversary's beliefs, both about themselves and about us, on their goals and on certain high level actions associated with those goals
3. Actions: A probabilistic network, representing the detailed relationships between adversary goals and the actions they are likely to perform to realize those goals

The goal component captures *what* the adversary is doing, the action component captures *how* the adversary might do it, and the rationale component infers *why* the individual is doing it. Due to the inherent uncertainty involved in adversary course of action prediction, we use Bayesian networks (Pearl 1988) as the main knowledge representation for the rationale and action networks. Each random variable (RV) involved in the Bayesian networks is classified into one of four classes: axioms, beliefs, goals and actions. Each RV class is described below:

- a Adversary axioms—represents the underlying beliefs of the adversary about themselves (vs. beliefs about our forces). This can range from an adversary's beliefs about his or her own capabilities to modeling a fanatic's belief of invulnerability. Axioms typically serve as inputs or explanations to the other RVs such as adversary goals
- b Adversary beliefs—represents the adversary's beliefs regarding our forces (e.g., an adversary may believe that the United States is on a crusade against them or that the United States is not carpet-bombing territory)
- c Adversary goals—represents the goals or desired end-states of the adversary. These goals are defined as either short-term or long-term in a goals list. Further we parti-

tion goals into two types: abstract and concrete. Abstract goals are those that cannot be executed (e.g., preserving launchers, damage US world opinion, defeating US foreign policy). They are satisfied by other abstract goals and also by concrete goals. Concrete goals are executable goals (e.g., repositioning launchers, contacting ambassadors, storing military equipment in civilian structures). Concrete goals can only be satisfied by concrete goals

d Adversary actions – represents the actions of the adversary that can typically be observed by friendly forces.

Figure 2 also shows feedback and explanation paths within the adversary intent inference (AII) model. Feedback from a human analyst, although unlikely to be totally certain, can be extremely valuable to the AII model, correcting and extending its intent inferencing logic. Explanation capabilities are essential in order for intelligence analysts, using AII, to understand why the AII model has reached particular inferences. The analysts must be able to inspect the reasoning paths used by AII so that they can develop a level of confidence in the output of the AII model.

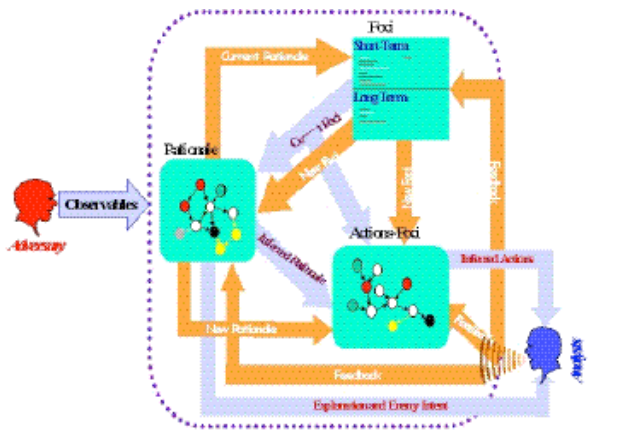


Figure 2. Adversary Intent Inference Model.

Collection and Production of Observables. There are two primary inputs to the adversary intent inferencing (AII) model—intelligence, surveillance, reconnaissance (ISR) data (observations of the actions of the adversary, collectively the *observables*) and direct analyst input. Unlike the majority of observables used to infer user intent, observables in the adversary intent domain cannot be pulled directly from an application. Instead, they must be discovered within the flow of information available to the decision makers at the strategic, operational and tactical levels of operations. An example of an observable might be that “the adversary is repositioning its SCUD launchers” or “an enemy tank unit is approaching one of the blue force’s key logistics and supply bases”. A large portion of these observables can be found in existing intelligence databases. Other observables may come from situation reports from fielded units or from online fact-sheets for the region of interest or from news reports. Observables may be gathered through friendly force sensor systems. We use the phrase

“sensor system” here to mean any system capable of collecting data on the adversary. Friendly force human intelligence operatives may also be considered to be sensor systems.

In collaboration with our ongoing research, our research partner, Lockheed Martin Advanced Technology Laboratories (ATL), is developing, under internal research and development funding, agent-based technology to collect sensor reports from three broad categories—tactical battlefield sensors, structured intelligence databases, unstructured data—process these reports and combine them to produce evidence to support higher level reasoning, and in particular to provide inputs to the adversary intent inferencing (AII) model. This technology is referred to as the Smart Agent Generation Engine or SAGE and is based on ATL’s Extendable Mobile Agent Architecture (EMAA) developed over the last seven years with DARPA and internal ATL funding (Whitebread and Jameson 1995; Hofmann et al. 2000; and McGrath, Chacon, and Whitebread 2000).

SAGE begins by analyzing the current evidence requirements of the AII model. By reviewing the state of the AII, SAGE attempts to prioritize evidence collection based on an analysis of which evidence will most aid with the disambiguation of the adversary’s intent. Once SAGE has this prioritized list of evidence requirements, it begins to decompose each evidence requirement into a sequenced set of data retrieval tasks. For example in order to determine whether the adversary is approaching our supply base, the velocity of the adversary relative to the supply base must be determined. The data retrieval tasks are collected into itineraries for software agents. These agents are then generated and dispatched to the appropriate distributed data sources. These data sources may be the high refresh rate outputs from level 1 fusion, the structured Joint Common Database (JCDB) or Military Intelligence Database (MIDB) or the unstructured news feeds. In the case of the unstructured data sources, the agents will likely request specific searches from evidence extraction and link discovery services.

Once the agents have collected the appropriate data, SAGE takes the data and combines it into evidence. This evidence combination process varies greatly in complexity. In order to reason about the range of a SCUD, the agent may just have had to return the value of the range field for the SCUD from the JCDB. In order to reason about whether adversary tanks are approaching a blue force supply base, historical values of the tanks’ position and velocities will be required as will the position of the supply base and any possible avenues of attack the adversary could follow. A further complexity is that evidence returned to AII by SAGE must have an associated probability value ranging from 0 to 1. Thus, SAGE must also estimate the probabilities associated with each piece of evidence.

The second input to the adversary intent inferencing (AII) model is analyst feedback. Feedback plays a critical role and effectively updating the intent model. Feedback from the analysts in adversarial intent must be inherently uncertain. This adds an additional level of criticality to the

explanation component. In particular, the intent model manages and maintains significantly more knowledge concerning the adversary than can be cognitively handled by a human analyst. Thus, by providing an explanation and even an exploration facility to the human analysts, we “open” up the intent model for complete inspection by the analysts in as organized manner as possible. In essence, we leverage the uncertainties in the analysts’ inferences in order to better adapt the intent model to cover larger contingencies and increase robustness.

Conclusions

We have outlined our assessment of the best approach to addressing adversary intent inferencing based on current research and our own expertise. We are fully aware of the fact that adversary intent inferencing is a highly complex problem in which even experts do not agree on many of the fundamental issues. Currently, there are factors that we cannot concretely and precisely address but hope to do so as our project progresses. For example, we realize that with regard to observables, both user intent and adversary intent domains must determine what types / kinds of observables need to be captured for effective intent inferencing. In the user intent domain, we can assume that all observables are available and are precise. In the adversary intent domain, however, observables may not be completely obtainable or even reliable due the fog and friction of war and to deception and subterfuge on the part on the adversary. Our work is in the definition phase and this paper reflects both our past experience and current plans. In the near future we will engage with various end-user groups in the Air Force intelligence community to build a preliminary adversary intent model and to identify the information sources to be accessed. We will develop a preliminary prototype of the adversary intent inferencing model by the end of 2002.

Acknowledgements

The research presented in this paper would not be possible without the involvement of our research partners, Lockheed Martin Advance Technology Laboratories and specifically Sergio Gigli and Axel Anaruk. We thank AFRL’s Information Institute for providing guidance and subject matter expertise and to AFRL’s Human Effectiveness Directorate (2d Lt Sabina Noll) for providing a workshop forum for discussing the theory and requirements for adversary decision modeling.

References

- Behler, R. Maj Gen 2001, “Homeland Information: AOC Can Coordinate U.S. Terror Defense,” *Defense News*, 13.
- Bell, B., Franke, J., and Mendenhall, H. 2000, “Leveraging Task Models for Team Intent Inference,” *Proceedings of the International Conference on Artificial Intelligence*.
- Bell, B., Santos Jr., E., and Brown, S. M. 2002, “Making Adversary Decision Modeling Tractable with Intent Inference and Information Fusion”, *Proceedings of the 11th Annual Computer Generated Force and Behavioral Representation Conference*.
- Fayette, D. F. 2001, “Effects-Based Operations: Application of new concepts, tactics, and software tools support the Air Force vision for effects-based operations”, *Air Force Research Laboratory Technology Horizons*, IF-00-15.
- Franke, J., Brown, S. M., Bell, B., and Mendenhall, H. 2000, “Enhancing Teamwork Through Team-Level Intent Inference,” *Proceedings of the International Conference on Artificial Intelligence*.
- Geddes, N. 1986, “The Use of Individual Differences in Inferring Human Operator Intentions,” *Proceedings of the Second Annual Aerospace Applications of Artificial Intelligence Conference*.
- Hofmann, M., Chacon, D., Mayer, G., and Whitebread, K. 2001, “CAST Agents: Network-Centric Fires Unleashed.” *2001 National Fire Control Symposium: Session: Automation of the Kill Chain*, Lihue, Hawaii.
- Joint Publication 3-13 1998, *Joint Doctrine for Information Operations*, Department of Defense.
- Joint Publication 2-01.3 2000, *Joint Tactics, Techniques, and Procedures for Joint Intelligence Preparation of the Battlespace*, Department of Defense.
- Llinas, J., Drury, C., Bialas, W., and Chen, A.C. 1998, “Studies and Analyses of Vulnerabilities in Aided Adversarial Decision Making,” AFRL-HE-WP-TR-1998-0099.
- Llinas, J., Drury, C., Jian, J. Y., Bisantz, A., and Younho Seong, Y. 1999, “Studies and Analyses of Aided Adversarial Decision Making Phase 2: Research on Human Trust in Automation,” AFRL-HE-WP-TR-1999-0216.
- McCrabb, M., Concept of Operations for Effects-Based Operations 2000, Draft paper for AFRL/IFTB, Version 2.0.
- McGrath, S., Chacon, D., and Whitebread, K. 2000, “Intelligent Mobile Agents in the Military Domain.” *Proceedings of the Fourth International Conference on Autonomous Agents 2000*.
- Pearl, J. 1988, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*, Morgan Kaufmann.
- Saba, M.G., and Santos Jr., E. 2000, “The Multi-Agent Distributed Goal Satisfaction System,” *Proceedings of the International ICSC Symposium on Multi-Agents and Mobile Agents in Virtual Organizations and E-Commerce (MAMA 2000)*, 389-394, Wollongong, Australia.
- See, J. and Kuperman, G. 1997, “Information Warfare: Evaluation of Operator Information Processing Models,” AFRL-HE-WP-TR-1997-0166.
- Whitebread, K. and Jameson, S. 1995, “Information Discovery in High-Volume, Frequently Changing Data,” *IEEE Expert Journal – Intelligent Systems and Applications*.