

A Cognitive Architecture for Adversary Intent Inferencing: Structure of Knowledge and Computation

Eugene Santos Jr.

University of Connecticut, Dept of Computer Science & Engineering, Storrs, CT 06269
eugene@cse.uconn.edu

ABSTRACT

Existing target-based and objectives-based (“strategy-to-task”) approaches to mission planning do not explicitly address the adversary’s decision-making processes. Obviously, the adversary’s courses of action (COA) are influenced in a cause-and-effect manner by actions taken by friendly forces. Given the iterative/interleaved nature of actions taken by enemy and friendly forces, mission planning must clearly take adversarial decision making into account especially during concurrent mission planning and execution. Currently, adversarial behaviour with regards to cause-and-effect are difficult to account for within the framework of existing planning approaches. This paper describes a cognitive architecture for computationally modeling, predicting, and explaining adversarial behaviors and COAs and proposes an integrated framework for mission planning. Our framework fits naturally within the Effects-Based Operations (EBO) approach to mission planning.

Keywords: Adversary Behavior, Inferred Behavior, Course of Action Prediction, Course of Action Analysis, Cognitive Architecture, Bayesian Networks, Reasoning, Computation, Enemy Course of Action

1. INTRODUCTION

Currently, target-based and objectives-based (“strategy-to-task”) approaches to planning do not explicitly address the adversary’s decision-making processes [6]. In particular, the adversary’s behaviors and courses of action (COA) are obviously influenced in a cause-and-effect manner which are actually difficult to account for within the framework of existing planning approaches. Furthermore, friendly COAs directly impact future adversary actions. Thus, to properly take into account adversarial decision-making and behavior in mission planning requires the capability to model adversaries in order to (1) predict enemy COAs, (2) infer adversarial intent, and (3) explain the rationale behind adversarial behaviour and intent. Such a model must capture a variety of attributes including effects of friendly and enemy behavior/action over time with regards to the current situation as well as soft factors such as enemy will, political system, and system of beliefs.

The foundations supporting the development of such adversary-aware mission planning systems are 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 [17]. 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 [13].” 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 [17]. 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 between them), we believe that a

necessary starting point is to model an adversary commander's intent. Intent inference involves deducing an individual's goals based on observations of that individual's actions [15]. 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 team intent modeling to address the general problem of the “enemy -as-a-system.”

Our goal in this paper is to present a cognitive architecture for adversarial modelling and provide a computational framework for adversary intent inference.

2. BACKGROUND AND APPROACH

Intent inference involves deducing an individual's goals based on observations of that individual's actions [15]. 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.

There are three major functions of intent inference [14]: *Descriptive* intent inference provides insight into the motivations behind actions that have just occurred. *Predictive* intent inference can anticipate future actions given the individual's inferred goals. *Diagnostic* intent inference arises from a targeted combination of predictive and descriptive models, which compares previous predictions against current knowledge. This works to reveal discrepancies in either the models (supporting a model adaptation and learning function) or the real world (identifying mistakes made by the individual).

Each function of intent inference can be dissected into three informational components [23,24,25]: The first, *interests and focus*, captures at a high level the direction of the individual's attention. The second, *actions and preferences*, describes the activities that can be used to carry out the goals that currently hold the individual's attention, with a focus on how the individual tends to carry them out. The third, *knowledge and reasoning*, provides insight into the deeper motivations behind the goals upon which the individual is focused and illuminates connections between goals. In other words, the first component captures what the individual is doing, the second captures how the individual might do it, and the third infers why the individual is doing it.

Applying the principles of modeling the what/how/why of individual intent, we see that our approach naturally integrates into the major themes of EBO. While the field of individual intent inference has historically focused on better improving the human-system interface, we contend that there is a natural isomorphism between our own prior work in the field of user and team intent inference [2,3,4,5,14,22,23,24,25] and the domain of adversary intent inference [1,6,29]. While the operational world surrounding an intent inference application would be very different, the inner mechanisms of intent inference map directly between domains.

While observables in the user intent domain stem from data collected from human use of systems, observables in the adversary intent domain take the form of tactical information derived from intelligence databases, observations of the tactical environment, and input from online human experts. In place of window events, keystrokes, and mouse movements, our system in the adversary intent domain uses information about adversary location, movements, and activities to drive its inference. 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.

Likewise, tactical goals will replace computer operational goals in the results of our intent inference. Descriptive intent inference in this case would result in identification of an adversary force's objectives and, given models of tactical

reasoning, could recommend appropriate reactions. Predictive intent inference would indicate expected activity by the adversary and explain the reasons behind that activity. Diagnostic intent inference could produce alerts of attempted subterfuge or uncover missteps on the part of the adversary.

Similarly, intent inference of echelons of adversary forces provides advantages reminiscent of those that arise from team intent inference. Identification of the goals of one adversary group can be used as a discriminator in identifying the goals of other subsets of the adversary. Intent inference across groups of the adversary could also result in the discovery of breakdowns among those groups, knowledge that can be used to our tactical advantage. We now present the cognitive architecture that forms the basis for our adversarial modeling.

3. ADVERSARY INTENT INFERENCE MODEL

To achieve adversarial intent inferencing requires the ability to (1) fuse information (observables) from sensors and intelligence sources regarding the adversary, (2) infer adversary intent and goals, and (3) predict adversary courses of action (COA). In total, adversary intent inferencing (AII) provides these three key functions while also taking into consideration a number of utility issues:

- AII must be able to explain the basis of its predictions; why is the adversary pursuing a predicted goal? What is driving the adversary to pursue these COAs? Must be able to model and take into account many factors including soft factors such as political environment, personality issues, adversarial religious beliefs, etc.
- AII must be able to adapt predictions based on history of events and observed enemy operations.
- AII must be dynamic and able to learn changes in the adversary behaviour and ultimately model pop-up adversaries; AII must be able to provide the capability for standing up models of new adversaries while avoiding the knowledge/information engineering bottleneck.

While the ultimate role and capabilities of AII still requires a great deal of long term research, from our first steps thus far, we can carefully identify capabilities that are likely achievable in the short-term for mission planning and wargaming to support Effects-based Operations (EBO), Predictive Battlespace Awareness (PBA), and Intelligent Preparation of the Battlespace (IPB).

As we discussed above, we derive our adversarial architecture from the formative components found in our user modeling approach to intent inferencing. In particular, we preserve the structure of the what/how/why model in order to provide a natural and intuitive decomposition of both the adversarial decision-making process and central knowledge-base. The benefits of such a decomposition are two-fold: First and foremost is the classic bottleneck of knowledge acquisition. Our decomposition provides a critical organizational structure in order to better capture/construct adversarial knowledge-bases/models in a manageable fashion. Secondly, with more modular components, this eases the issues of computational complexity and validation/auditability of the inferencing process.

The components of our adversary intent inferencing model, and the interactions between these components, are shown in Figure 1. The three core components that comprise our architecture and functions are as follows:

- **Goals:** Probabilistically prioritized short- and long-term goals list, representing adversary intents, objectives or foci
- **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
- **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 [18] 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 (X)** – 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 (B)** – 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 (G)** – 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 partition 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).
- (d) **Adversary actions (A)** – represents the actions of the adversary that can typically be observed by friendly forces.

These four random variable types are arranged in the two networks: rationale network and action network. The rationale network contains all of the Belief (B), Axiom(X), and Goal (G) variables, as well as any Action (A) variables which have goals as inputs. This network is used to infer what short and long term goals the adversary may have. Once the goals are determined, the action network is used to reason on what the most likely actions will be that the adversary may carry out. The action net contains the entire set of Action (A) variables and any Goal (G) variables that could be considered as actions. As a rule, Belief variables are independent and serve as inputs to Axioms or Goals. Axioms have Beliefs as inputs and serve as inputs to Goals and other Axioms. Goals have Axioms and Beliefs as inputs and serve as inputs to Actions or other Goals. Actions have only Goals as inputs and can only be inputs to other Actions. There are additional rules for structuring the relationships between random variables that are not directly described here due to space limitations but can be seen in the example networks. Figure 2 depicts a rationale network and an action network.

The AII process (as shown in Figure 1) works iteratively as follows:

1. Observables regarding the adversary such as actions and beliefs are set as evidence in both rationale and action networks (depicted as red nodes in figure). Also, feedback from analyst is set as evidence.
2. Current short- and long-term enemy foci from the foci lists are also set as evidence in both networks (depicted as green nodes).
3. The rationale network is then used to infer new goals which are set as evidence for the action network.
4. The action network is now used to predict adversarial actions.
5. The analyst is presented with the inferred goals and predicted actions.
6. The analyst provides feedback in terms of corrected goals and actions if desired.
7. The goals list is updated based on newly inferred goals and current strength of existing goals. If goals exceed a given threshold value, they are added to the list. If goals fall below a set threshold, they are removed. If goals in the short-term list persist beyond a given time threshold, they become long-term goals.
8. Go to step 1.

Due to space considerations for this paper, we have omitted details regarding specific functions/formulas used. The inference process on both the rationale and action networks is based on belief updating [18]. In essence, given a target random variable R and evidence set E, belief updating computes $P(R|E)$ assuming that random variables have two states (true/false) for simplicity of discussion. In the following section, we give a brief overview of Bayesian networks in order to provide a better understanding of uncertainty modeling issues involved in our framework.

As we can see in the above process, the adversary model is capable of adapting to changes in the adversaries goals and intentions over time as reflected in the enemy foci lists. Also note that there are 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.

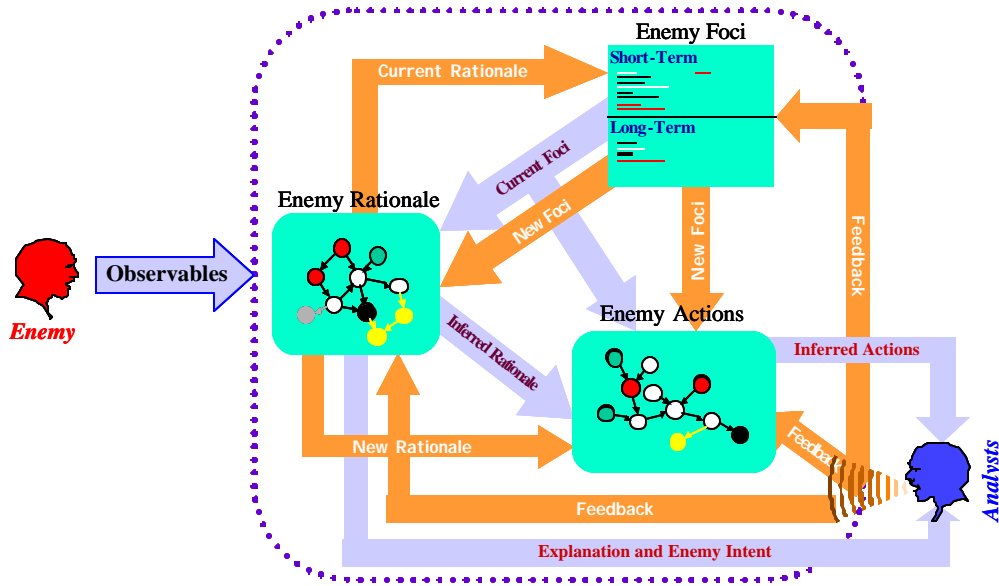


Figure 1. AII Process

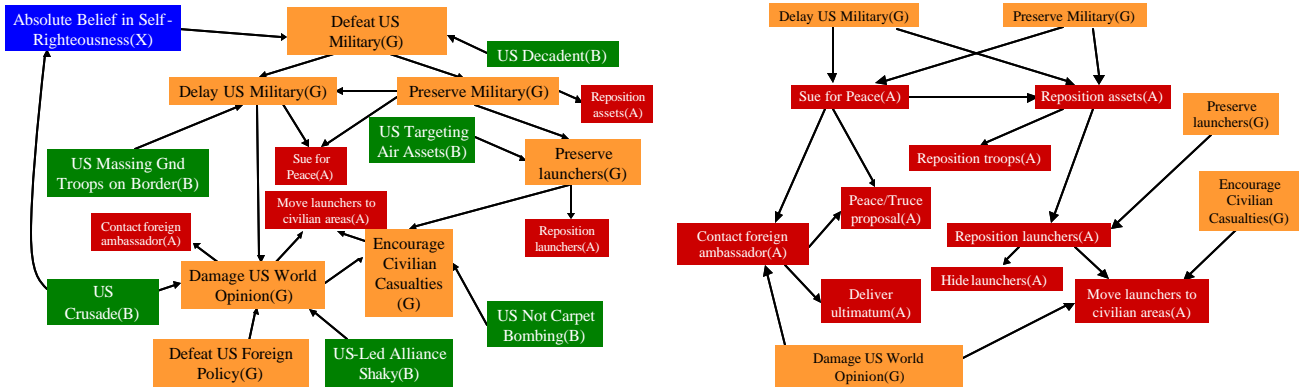


Figure 2. Rationale Network (left) and Action Network (right). The random variables are labeled according to their categories.

4. UNCERTAINTY

Our rationale and action networks must capture the uncertainties inherent in the adversarial model as well as the uncertainties found in the observables. In probabilistic reasoning, random variables (abbreviated, r.v.) are used to represent events and/or objects in the world. By making various instantiations to these r.v.s, we can model the current state of the world probabilistically. Thus, this will involve computing joint probabilities of the given r.v.s. Unfortunately, the task is nearly impossible without additional information concerning relationships between the r.v.s. In the worst case, we would need the probabilities of every instantiation combination which is combinatorially explosive. On the other hand, consider the chain rule as follows:

$$P(A_1, A_2, A_3, A_4, A_5) = P(A_1 | A_2, A_3, A_4, A_5) P(A_2 | A_3, A_4, A_5) P(A_3 | A_4, A_5) P(A_4 | A_5) P(A_5).$$

Bayesian networks [18] take this process further by making the important observation that certain r.v. pairs may become uncorrelated once information concerning some other r.v.(s) is known. More precisely, we may have the following independence condition:

$$P(A | C_1, \dots, C_n, U) = P(A | C_1, \dots, C_n)$$

for some collection of r.v.s U . Intuitively, we can interpret this as saying that A is determined by C_1, \dots, C_n regardless of U .

Combined with the chain rule, these conditional independencies allow us to replace the terms in the chain rule with the smaller conditionals. Thus, instead of explicitly keeping the joint probabilities, all we need are smaller conditional probability tables which we can then use to compute the joint probabilities.

In Bayesian networks, these conditional dependencies are represented as a directed acyclic graph of r.v. relationships. Directed arcs between r.v.s represent direct conditional dependencies. When all the parents of a given r.v. A are instantiated, that r.v. is said to be conditionally independent of the remaining r.v.s which are not descendants of A given its parents. (For more details on this, see d-separation in [18].)

For example, let's consider the following story: Mary walks outside and finds that the street and lawn are wet. She concludes that it has just rained recently. Furthermore, she decides that she does not need to water her climbing roses. Assume that Mary used the following set of rules:

```
rain or sprinklers --> street = wet
rain or sprinklers --> lawn = wet
lawn = wet --> soil = moist
soil = moist --> roses = okay
```

We can directly transform these into a graph. Now, by considering each variable as a r.v. with possible states of {true,false}, we can construct conditional probability tables for r.v. which reflects our knowledge of the world. The joint probability of the world where the roses are okay, the soil is dry, the lawn is wet, the street is wet, the sprinklers are off and it is raining is computed as follows:

$$P(\text{sprinklers} = F, \text{rain} = T, \text{street} = \text{wet}, \text{lawn} = \text{wet}, \text{soil} = \text{dry}, \text{roses} = \text{okay}) = P(\text{roses} = \text{okay} \mid \text{soil} = \text{dry}) * P(\text{soil} = \text{dry} \mid \text{lawn} = \text{wet}) * P(\text{lawn} = \text{wet} \mid \text{rain} = T, \text{sprinklers} = F) * P(\text{street} = \text{wet} \mid \text{rain} = T, \text{sprinklers} = F) * P(\text{sprinklers} = F) * P(\text{rain} = T)$$

Substituting the appropriate numbers from the tables, we get $0.2 * 0.1 * 1.0 * 1.0 * 0.6 * 0.7 = 0.0084$ as the probability of this scenario.

There are two types of computations performed with Bayesian Networks: belief updating and belief revision [18]. Belief updating concerns the computation of probabilities over random variables, while belief revision concerns finding the maximally probable global assignment.

Belief revision can be used for modeling explanatory/diagnostic tasks. Basically, some evidence or observation is given to us, and our task is to come up with a set of hypothesis that together constitute the most satisfactory explanation/interpretation of the evidence at hand. This process has also been considered abductive reasoning in one form or another [19]. More formally, if W is the set of all r.v.s in our given Bayesian network and e is our given evidence, i.e., e represents a set of instantiations made on a subset of W , any complete instantiations to all the r.v.s in W which is consistent with e will be called an explanation or interpretation of e . Our problem is to find an explanation w^* such that

$$P(w^* \mid e) = \max P(w \mid e).$$

w^* is called the "most-probable explanation." Note that to compute the most-probable explanation for e , it is sufficient to determine the complete assignment consistent with e whose joint probability is maximal. In this case, $P(e)$ is simple a constant factor. Intuitively, we can think of the non-evidence r.v.s in W as possible hypotheses for e .

Belief updating on the other hand is interested only in the marginal probabilities of a subset of r.v.s given the evidence. Typically, it is to determine the best instantiation of a single r.v. given the evidence. For example, let the evidence e be

the observation that the roses are okay and the condition of our lawn be our focus. Our goal is to now determine the probability that our lawn is either wet or dry given the observation. The solution then becomes –

$$\begin{aligned}P(\text{lawn} = \text{dry} \mid \text{roses} = \text{okay}) &= 0.1190 \\P(\text{lawn} = \text{wet} \mid \text{roses} = \text{okay}) &= 0.8810\end{aligned}$$

Although performing belief revision and updating (even approximating methods) have been shown to be NP-hard, there exist special network topologies for which certain algorithms perform well such as polytrees [18]. Various approaches to reasoning with Bayesian Networks include A* search, stochastic simulation, integer programming, and message passing [20,26,28,21,27,33].

5. WHAT'S NEXT?

Currently, the AII adapts by capturing temporal changes in adversarial activities through the short-term and long-term foci lists. While we believe that this is one of the most critical capabilities that must be provided for useful adversarial behavior prediction, additional adaptive capabilities are needed to ultimately solve problems such as pop-up adversaries. In the worst case, one of the primary difficulties with pop-up adversaries is the potential lack of knowledge or incompleteness of information available to build our networks apriori. To address this challenge, the AII must be capable of automatically updating its network models by adding or removing nodes as observations, predictions, and feedback is garnered regarding the adversary. In Figure 1, such changes are made in the feedback stage where the yellow nodes represent new additions and the gray node represents deletion. Our vision initially for achieving this adaptability focuses on two elements: (1) the identification of a need for new knowledge or removal of old/incorrect knowledge, and (2) the construction/destruction of knowledge. For (1), we can detect this situation when the recent feedbacks from the analyst are significantly different from or contradicts the AII predictions. In this case, the answers provided by AII are inconsistent either because of *incompleteness* -- insufficient information (relevant nodes in the networks) currently captured, or *incorrectness* -- there is information that is incorrectly captured. To address incompleteness, we envision a library of knowledge nuggets which are small network fragments corresponding to simple rules or templates. When additional information is needed, the library is referenced and the appropriate knowledge nuggets are obtained and introduced (like the yellow nodes) into the networks. For incorrectness, we can initiate sensitivity analysis on the current networks to identify the nodes that are the significant cause of the inconsistency. Once identified, they and nearby nodes are removed and stored back into the library as knowledge nuggets. One of the side benefits of this approach is that adversarial models can be constructed on the fly and computational costs from inferencing can be better controlled. Also, the knowledge nuggets should be easier to formulate and validated while being applicable to various large collections of different adversaries.

Clearly, what we have just outlined so far in this paper reflects our current beliefs on the best approach to addressing adversary intent inference based on current research and our own expertise. We are quite aware of the fact that adversary intent inference is a highly complex problem in which even experts do not agree on many of the fundamental salient points. Currently, there are factors that we cannot concretely and precisely address but hope to do so as our efforts during the project provide us with more insights both from successes and failures. 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 inference. However, for user intent, we can assume that all observables are available and are precise. For adversary intent, observables may not be completely obtainable or even reliable due the factors arising from the fog and friction of war effects to deception and subterfuge on the part on the adversary. In the next section, we present a longer term vision for properly integrating adversary intent inferencing that we believe will help address many of these issues.

6. THE BIG PICTURE

To recap, understanding the adversary is a long and well-known fundamental need for effective military planning and operations. A major challenge we face today is in providing the ability to explain as well as predict enemy intentions, goals, behaviours, plans, and actions and then effectively integrating this information into blue forces mission planning and execution. This is further complicated by the online nature of real-world operations in which actions taken by the

blue forces will undoubtedly affect future actions and goals of the adversary. Also, the inherent uncertainty such as the fog-of-war and enemy deception constrains the information obtainable both in terms of quality and temporal availability. All of this must also take into consideration the limited resources available for information gathering. In this section, we present some initial thoughts and ideas on how-to and what it takes to properly account for the adversary.¹

We propose a model and architecture for adversary intent inferencing and course-of-action prediction with dynamic information fusion and gathering. Our goal is to provide a unified approach that is composed of 3 major elements: (1) adversary modelling and intent inferencing, (2) adversary plan recognition and course-of-action prediction, and (3) adaptive information tuning and fusion. The first component provides the basic capability to model the adversary and explain/predict their behaviour from observations gathered by component (3). The plan recognizer and adversary COA prediction then uses the predicted behaviour from (1) and observables from (3). The results of (2) can then be directly used by blue planning systems. Our information tuning/gathering system then uses the results from (1) and (2) to retask themselves to either validate a prior observable or search for new observables that can further improve the predictions of (1) and (2). This is an iterative/online process that we will now consider in more detail.

Figure 3 presents a high-level view of the proposed Adversary Intent/COA recognizing system proposed. As shown, there are three key components: Adversary intent inferencing, case-based plan recognition, and the adaptive information system. We now present these three components in detail.

Adversary Intent Inferencing. The adversary intent inferencer is responsible for determining red goals including strategic and high-level tactical goals and actions.

The AII takes input from 3 sources:

1. Evidence/observables from AIS – these may be observations straight from the battlefield, recon, sensor arrays, etc.
2. Projected Red COAs from plan recognizer – these are critical to better guiding the AII in identifying red goals allowing for more flexibility in dealing especially with soft factors without the explosion of uncertainty.
3. Analyst input – critical to merging and working with human analysts.

A complete system must encompass all 3 in order to complete the cycle of red forces analyses and for proper incorporation into blue forces planning. With regards to input 2, hard factors as opposed to soft factors can be characterized as factors that are measurable, observable, or physical in nature. Soft factors, on the other hand, include intangibles such as enemy will, political influences, personality, human behaviour, and fundamental belief systems.

Case-Based Plan Recognition. The plan recognizer is responsible for providing a key piece of information the analyst desires. That is it takes as input targeted information from the AIS (e.g., blue goals and plans along with red actions) and high-level red goals predicted from AII, and it produces as output possible red courses of actions (COAs). This output will be directed to the tuner mechanism of AIS, to AII as feedback, and to the analyst. Analysts can then make decisions (including subsequent parameter tuning of the AIS) based upon the potential red COAs.

To perform such predictions the system must reason about information concerning both red and blue, because possible responses from each will interact. One obvious interaction is the relative physical deployment locations of forces. The COAs predict employment of red forces, but this prediction also depends on current blue deployment and what red knows of it. Managing these interactions well crucially depends upon knowing both blue goals (given from AIS) and likely red goals (inferred by AII). Knowledge about goal interactions (e.g., goal conflict. See [31].) and the possession of a general theory of goal change and management [7,8,10,32,12] enables the recognizer to constrain the possible high level interpretations of events using the Meta-AQUA system [9]. Statistical methods used by the incremental case-based plan recognition system [16] are thus more tractable.

The COAs are also just a prediction, so the recognition component will directly provide AIS with future events to monitor in order to strengthen or modify such predictions. These time sensitive information requests are very similar to rationale-based planning monitors [30], but they are spawned in response to inferred red plans rather than known blue

¹ The ideas presented here resulted from joint work between the author, Scott Deloach, and Michael Cox.

plans. Finally, particular blue actions interact with other blue intentions. Therefore AIS will be directed to monitor a select key set of future blue events and report to the plan recognition component.

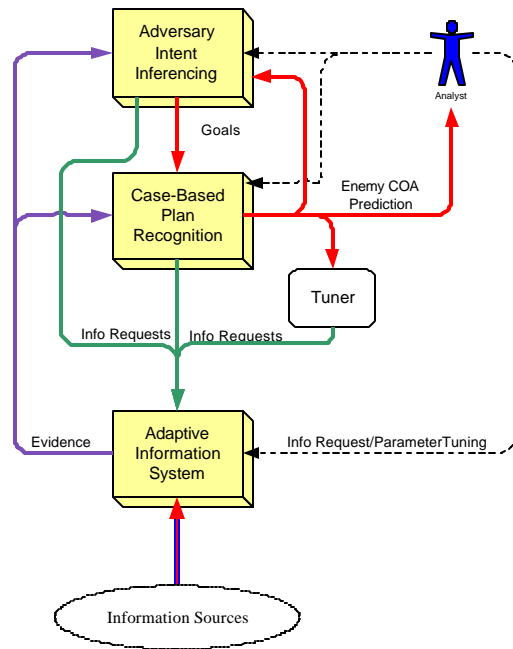


Figure 3. Big Picture

To summarize, the functionality of the hybrid case-based plan recognition system is as follows.

1. To produce predictions of enemy COAs based on predicted red and known blue intentions.
2. To anticipate possible interactions of red and blue given current information by requesting states to be monitored by AIS. The anticipations may be both positive and negative. The former signals hypothesized opportunities, whereas the negative signals warnings.
3. To update predictions and information requests as new evidence arrives and as inferred red or blue goals change.
4. To produce warnings to analysts when actions on one blue component interferes with the commitments of another blue component.

Adaptive Information System (AIS). The AIS is responsible for providing information from various sources to provide evidence for the other two components. The information provided is based on analyst input, requests for further information from the Plan Recognizer, AII, or new information requests based on predicted enemy courses of action. Requests for further evidence may be requests for current or past evidence, or they may be requests for future information as it becomes available. The AIS is tunable based on time or quality. The AIS can be asked to provide the best information possible by a certain deadline or to provide information at a specified confidence level. The AIS consists of an organization of intelligent information agents that reorganize to provide the required information while reducing communication and processing overhead. The individual information agents are relatively simple and only know how to produce certain pieces of information. This information can be gathered directly from known information sources or by combining outputs from known sources or other information agents. The AIS adapts to changes in information requests by modifying its organization to provide the right information, at the right time, at the right quality. Adding new intelligent information agents and updating the team's organizational knowledge [11] can extend the AIS incrementally during execution. This capability allows the AIS to automatically adapt quickly to changing requirements as well as to long term changes in the information environment without having to "rebuild" the entire system.

Tuner. The fourth component shown Figure 3 is the AIS “Tuner”, which is capable of automatically modifying the AIS information requests based on the enemy courses of action predicted by the Plan Recognizer. This component is not integral to the overall system and may not be included in the initial versions of the system.

Summary. We have presented a unified approach to incorporating the adversary into an online system for explanation and prediction of adversary behaviours and actions. It takes into account adversary reactions to blue force actions and addresses the need to retask information gathering resources to better tune adversary predictions. We presented an architecture that can be integrated into blue forces mission planning and execution.

7. PROJECT STATUS AND RELATED EFFORTS

In this paper, we have presented a cognitive architecture for adversarial intent inferencing for use in future mission planning systems. The Adversarial Intent Inference for Predictive Battlespace Awareness Project is a basic 6.1/6.2 research project sponsored by the AFRL/IF’s Information Institute Research Program. Our goal has been to design and develop advanced adversarial modeling and prediction tools that can provide the necessary enabling technologies to support AF mission planning and execution as well as wargaming needs such as EBO, Predictive Battlespace Awareness (PBA), Intelligent Preparation of the Battlespace (IPB), Information Fusion, etc. Currently, we have built a prototype system to model the adversary’s behavior based on probabilistic models (Bayesian Networks/Influence Diagrams) and evaluated its effectiveness as a proof-of-concept. Our prototype simulates the “Battle at Khafji” scenario during the Persian Gulf War by predicting and updating the predicted actions of the adversary over time as the events unfolded.

The architecture has also been inserted into military wargaming environments [29]. In the current world environment, the rapidly changing dynamics of organizational adversaries are increasing the difficulty for Military analysts and planners to accurately predict potential actions. As an integral part of the planning process we need to assess our planning strategies against the range of potential adversarial actions. This research project investigates the feasibility of utilizing AII as a core element within a predictive simulation to establish emergent adversarial behavior. It is our desire to use this intelligent adversary to *generate alternative futures* in performing Course Of Action (COA) analysis. Such a system will allow planners to gauge and evaluate the effectiveness of alternative plans under varying actions and reactions.

8. ACKNOWLEDGEMENTS

The author would like to acknowledge the significant contributions to this project made by Sergio Gigli and Frank Vetesi (Lockheed Martin, Advanced Technology Laboratories), Robert Hillman (Air Force Research Laboratory, Information Directorate), Joshua Surman (University of Buffalo), and Ben Bell (CHI Systems). I would also like to especially thank John Graniero (Air Force Research Laboratory, Information Institute), Don Monk (Air Force Research Laboratory, Human Effectiveness), and Scott Brown (USAF) for their tremendous support towards establishing this research project. This work has been supported in part by a grant from the Air Force Research Labs, Information Directorate, Grant No.F30602-01-1-0595 through the Information Institute Research Initiative.

9. REFERENCES

1. Bell, B., Santos, E., Jr., and Brown, S. M., "Making Adversary Decision Modeling Tractable with Intent Inference and Information Fusion," *Proceedings of the 11th Conference on Computer Generated Forces and Behavioral Representation* (pp. 535-542), 2002.
2. Brown, S. M. and Santos, E., Jr., "Active User Interfaces", *IDIS TR No. 101*, Intelligent Distributed Information Systems Laboratory, University of Connecticut, 1999.
3. Brown, S.M., Santos, E., Jr., and Banks, S.B., "Utility Theory-Based User Models for Intelligent Interface Agents", *Proc. of the 12th Biennial Conf. of the Canadian Society for Computational Studies of Intelligence*, 379-393, 1998.
4. Brown, S.M., Santos, E., Jr., and Banks, S.B., "Active User Interfaces for Building Decision-Theoretic Systems," *Proceedings of the 1st Asia-Pacific Conference on Intelligent Agent Technology*, 244-253, Hong Kong, 1999.
5. Brown, Scott M., Santos, Eugene, Jr., Banks, Sheila B., and Oxley, Mark, "Using Explicit Requirements and Metrics for Interface Agent User Model Correction," *Proceedings of the Second International Conference on Autonomous Agents*, 1-7, Minneapolis/St. Paul, MN, 1998.
6. Brown, S. M., Santos, E., Jr., and Bell, B., "Knowledge Acquisition for Adversary Course of Action Prediction Models," *Proceedings of the AAAI Fall Symposium on Intent Inferencing for Users, Teams, and Adversaries*. North, 2002.
7. Cox, M. T., "A conflict of metaphors: Modeling the planning process," *Proceedings of 2000 Summer Computer Simulation Conference* (pp. 666-671), 2000.
8. Cox, M. T., Edwin, G., Balasubramanian, K., and Elahi, M., "Multiagent goal transformation and mixed-initiative planning using Prodigy/Agent," *Proceedings of the 5th World Multiconference on Systemics, Cybernetics and Informatics*, Vol. VII (pp. 1-6), 2001.
9. Cox, M. T. and Ram, A., "Introspective Multistrategy Learning: On the construction of learning strategies," *Artificial Intelligence*, **112**, 1-55, 1999.
10. Cox, M. T. and Veloso, M. M., "Goal Transformations in Continuous Planning," *Proceedings of the 1998 AAAI Fall Symposium on Distributed Continual Planning* (pp. 23-30), 1998.
11. DeLoach, S. A., "Modeling Organizational Rules in the Multiagent Systems Engineering Methodology," *Proceedings of the 15th Canadian Conference on Artificial Intelligence*, LNAI 2338 (pp. 1 – 15), 2002.
12. Edwin, G., "Comas: coordination in multiagent systems," Master's thesis, Wright State University, Dayton, OH, 2001.
13. Fayette, D. F., "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, 2001
14. Franke, J., Brown, S. M., Bell, B., and Mendenhall, H., "Enhancing Teamwork Through Team-Level Intent Inference," *Proceedings of the International Conference on Artificial Intelligence*, 2000.
15. Geddes, N., "The Use of Individual Differences in Inferring Human Operator Intentions," *Proceedings of the Second Annual Aerospace Applications of Artificial Intelligence Conference*, 1986.
16. Kerkez, B. and Cox, M. T., "Incremental Case-Based Plan Recognition Using State Indices," *Case-Based Reasoning Research and Development: Proceedings of the 4th International Conference on Case-Based Reasoning* (pp. 291-305), 2001.
17. McCrabb, M., "Concept of Operations for Effects-Based Operations 2000," Draft paper for AFRL/IFTB, Version 2.0.
18. Pearl, J., *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*, Morgan Kaufmann, 1988.

19. Santos, E., Jr., "A Linear Constraint Satisfaction Approach to Cost-Based Abduction," *Artificial Intelligence* 65(1), 1-28, 1994.
20. Santos, E., Jr., "On Linear Potential Functions for Approximating Bayesian Computations," *Journal of the ACM* 43(3), 399-430, 1996.
21. Santos, Eugene, Jr., "On Multiple Spline Approximations for Bayesian Computations," *Annals of Mathematics and Artificial Intelligence* 20(1-4), 267-300, 1997.
22. Santos, E., Jr., Brown, S.M., Lejter, M., Ngai, G., Banks, S.B., and Stytz, M.R., "Dynamic User Model Construction with Bayesian Networks for Intelligent Information Queries," *Proceedings of the 12th International FLAIRS Conference*, 3-7, 1999.
23. Santos, E., Jr., Brown, S. M., and Nguyen, H., "Medical Document Information Retrieval Through Active User Interfaces," *Proceedings of the 2000 International Conference on Artificial Intelligence (IC-AI '2000)*, Las Vegas, NV, 2000.
24. Santos, E., Jr., Nguyen, H., and Brown, S. M., "Kavanah: An Active User Interface Information Retrieval Application," *Proceedings of the 2nd Asia-Pacific Conference on Intelligent Agent Technology*, 412-423, 2001.
25. Santos, E., Jr., Nguyen, H., Zhao, Q., and Pukinskis, E., "Empirical Evaluation of Adaptive User Modeling in a Medical Information Retrieval Application," to appear in *Proceedings of the 9th International Conference on User Modeling*, Pittsburgh, PA, 2003.
26. Santos, E., Jr., and Shimony, S. E., "Deterministic Approximation of Marginal Probabilities in Bayes Nets," *IEEE Transactions on Systems, Man, and Cybernetics Part A: Systems and Humans* 28(4), 377-393, 1998.
27. Santos, E., Jr., Shimony, S. E., and Williams, E., "Hybrid Algorithms for Approximate Belief Updating in Bayes Nets," *International Journal of Approximate Reasoning* 17(2-3), 191-216, 1997.
28. Shimony, S. E. and Santos, E., Jr., "Exploiting Case-Based Independence for Approximating Marginal Probabilities," *International Journal of Approximate Reasoning* 14(1), 25-54, 1996.
29. Surman, J., Hillman, R., and Santos, E., Jr., "Adversarial Inferencing for Generating Dynamic Adversary Behavior," too appear in *Proceedings of the SPIE 17th Annual International Symposium on Aerospace/Defense Sensing and Controls: AeroSense 2003*, Orlando, FL, 2003
30. Veloso, M. M., Pollack, M. E., and Cox, M. T., "Rationale-based monitoring for continuous planning in dynamic environments," *Proceedings of the Fourth International Conference on Artificial Intelligence Planning Systems* (pp. 171-179), 1998.
31. Wilensky, R., *Planning and understanding: A computational approach to human reasoning*, Reading, MA: Addison-Wesley Publishing, 1983.
32. Zhang, C., "Cognitive models for mixed-initiative planning," Master's thesis, Wright State University, Dayton, OH, 2002.
33. Zhong, X. and Santos, E., Jr., "Directing Genetic Algorithms for Probabilistic Reasoning Through Reinforcement Learning," *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems* 8(2), 167-185, 2000.