Intelligent information assistance for coalition operations

Felipe Meneguzzi, Jean Oh, Katia Sycara
Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213, USA
{meneguzz|jeanoh|katia}@cs.cmu.edu

Abstract. The challenge of ensuring that the right information is available to the right personnel at the right time and in the right format is becoming increasingly important as military forces adopt a vision of *Network Centric Warfare and Operations*. This is particularly true in a dynamic environment with prevalent uncertainty, where various types and sources of information must be consulted while the information requirements change during planing and re-planning. The heterogeneity in doctrine and information systems inherent to coalition operations compounds this problem, particularly from the point of view of human users that must not only plan and execute missions but also keep in mind all potentially relevant sources of information, leading to cognitive overload. We have developed the ANTicipatory Information and Planning Agent (ANTIPA) [11] to manage information adaptively in order to mitigate user cognitive overload. To this end, the agent brings information to the user as a result of user requests, but, most crucially, it proactively predicts the user's prospective information needs by recognizing the user's plan; prefetches information that is likely to be used in the future; and offers the information when it is relevant to the current or future planning decisions. In this paper we review the ANTIPA architecture and describe potential applications of ANTIPA in coalition operations.

1 Introduction

Ensuring that the right information is available to the right personnel at the right time and in the right format is an often elusive task which is becoming increasingly important as military forces adopt a vision of *Network Centric Warfare and Operations*. This is particularly true in a dynamic environment with prevalent uncertainty, where people must consult various types of information in their decision-making processes while the information requirements rapidly change during planing and re-planning. As a result, users who must make time-critical decisions in information intensive tasks are cognitively overloaded by the planning activities and the information requirements of planning and re-planning. These difficulties are exacerbated by the heterogeneity in doctrine and information systems inherent to coalition operations.

For example, consider a military scenario where a coalition commander must plan (or re-plan) a critical mission in a fast-changing environment where information flows among elements of multiple military forces of varying nationalities. Due to uncertainty and dynamics in the environment, the commander must constantly collect up to date information to ensure the success of the mission; reason about the feasibility of the current plan; and synchronize with other involved commanders (so that the overall plan is coherent). Here, information that the commander must manage include intelligence reports, observations from the field, plan steps that must be executed, synchronization constraints, alerts, and others. Managing a multitude of information elements in a fast changing environment is a challenging task prone to failures and oversight, even if a commander is very familiar with the information sources from his/her own service. This challenge is compounded by the fact that during coalition operations there might be information sources that a commander is either not aware of or is not familiar with.

In this context, we develop an information agent that can manage information adaptively so that the users can focus on planning activities without being overwhelmed by information-gathering activities. Thus, we developed the *ANTicipatory Information and Planning Agent* (ANTIPA) [11] to address the informational needs of a human user.[1] In accomplishing its goal of helping the user to plan while minimizing distractions due to information needs, the agent proactively *predicts* the user's prospective information needs by recognizing the user's plan through monitoring its interaction with the system; *optimizes* information gathering; and *presents* information in a way that alleviates the user's cognitive load. Our current prototype uses a decision-theoretic approach for plan recognition, while other approaches use sequential decision-making models to design how an assistant agent should choose an optimal action based on its belief about the user's current state [6, 7]. Key to our approach is that the decision-theoretic model is used in the representation of how the user makes decisions, allowing the prediction of how this user will behave in the future. The predicted user plan provides a set of goals for the *ANTIPA* agent, for which the agent plans and executes assistive actions asynchronously. This separation allows the *ANTIPA* agent to have a far richer planning capability when compared to sequential action selections.

It is important to note that the goal of this research is not to guide the user in finding optimal planning solutions, but instead, the agent aims to optimize information management such that acquired information anticipates the user's information needs for planning decisions. As opposed to directing the user to make optimal decisions with respect to a certain objective (as in decision-support systems), we aim to design an agent that can maximize the support to help the user in making decisions based on her own criteria and judgement. From the user's perspective, independent decision making is crucial in many problem domains including military planning, educational support systems, and assistive living technologies for the disabled and the elderly.

In this paper, we describe potential applications of ANTIPA in coalition operations ranging from a commander in headquarters planning operations while trying to collate relevant information, to a commander on the ground receiving important information to the execution of the current part of its battle plan.

We start the paper by formalizing the problem our agent architecture was designed to address in Section 2. We then proceed to describing the ANTIPA agent architecture in Section 3, describing each of its components. With the basic principles of the architecture defined, we describe potential areas of application for the the ANTIPA architecture in Section 5, followed by related work that can be leveraged

[1] We shall use the term *user* throughout this paper to refer to the person utilizing ANTIPA.

2 Problem Definition

We define information-dependent planning problems as a class of planning problems where a user (or a planner) must access various types of information sources to acquire current information that is required for executing certain actions. Here, in addition to domain-specific planning objectives, the user must take the cost of getting information into consideration in selecting actions. Furthermore, the quality of information (that also depends on the source of information) affects the user's transition to another state after taking the action.

For instance, consider a student preparing for a final exam by reviewing selected topics covered in a semester. When a question is encountered, the student may search online for a quick answer by taking a risk that the answer found may be incorrect, or email the teacher and wait to get a generally more precise answer. The outcome of the student's *action* (to understand the concept) depends on the quality of information, which in turn depends on the source from which it has come. For example, by receiving high-quality information, the student's *state* regarding the understanding of a certain concept is more likely to *transition* from `not-learned` to `learned`, thus increasing the chance of getting a better grade (*reward*) in the final exam.

Given that the user is trying to solve an information-dependent planning problem, we design an agent that can adaptively identify and manage the user's information needs to facilitate the user's actions. The agent cannot directly observe the user's true states nor the actions that the user has taken; therefore, it must infer the user's state from a series of primitive sensory data known as *observations*. For instance, in the student example, possible observations include the keywords that the user types into search engines or a set of documents that the user opens.

3 The *ANTIPA* agent architecture

In this section we review the ANTicipatory Information and Planning Agent (*ANTIPA*) architecture [11], which we use to manage a user's information requirements by recognizing a users plan through observing her behavior. Figure 1 depicts the high-level architecture of an *ANTIPA* agent where the agent processes are contained within the dashed box; the cloud represents the agent's observations; rounded boxes represent data structures; and rectangle stacks represent reasoning tasks. The basic interactions between the user and the agent are: the agent observes some of the user's planning activities; and the agent may present information to the user. At deployment time, the agent is supplied with two inputs: a *domain description* representing the user's planning problem (*e.g.*, state-based planning problems, plan libraries, workflow, or similar activity representations); and an *information catalog* that describes a set of properties of information sources from which the agent can retrieve information. These two inputs to the information agent are shown as the two rectangles at the top left of Figure 1.

As part of the process of deciding on collecting and presenting information, the agent's reasoning process tries to accomplish four main objectives. First, the agent must identify the current state of the user from a sequence of observations on user activities. Second, the agent needs to predict the user's information

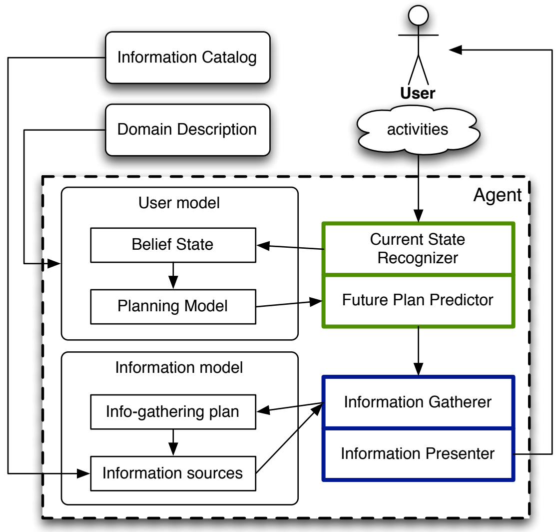

Figure 1: The *ANTIPA* agent architecture.

needs that are changing dynamically through time. In order to accomplish this, the agent needs to identify the user's high-level goals and a set of planned actions to achieve these goals, in a process known as plan recognition [2]. If the agent can recognize the user's goals and plans, then the agent can infer the information needs associated with these planning activities. Third, the agent needs to construct a plan for collecting the information from various information sources. This plan must consider the tradeoff between obtaining the information of which the user is likely to make the most use and satisfying temporal deadlines that certain information must be obtained before a specific time point to be useful. Finally, the agent must decide when to offer certain information to the user based on its belief about the user's current state. In order to accomplish these objectives, the *ANTIPA* architecture is composed of four main components as follows:

- current state recognizer;
- future plan predictor;
- information gatherer; and
- information presenter.

We shall review these components in the following sections.

3.1 Current State Recognizer

The agent models the user's current state (which the agent cannot observe directly) as a probability distribution over a set of possible states, known as a *belief state*. When the agent perceives a new observation from the user interacting with an environment, Current State Recognizer updates its belief state such that the updated belief state best explains the observations. The updated belief state then triggers other components to adjust accordingly (*e.g.*, the agent can determine a set of information to present immediately according to the current belief state).

3.2 Future Plan Predictor

Given a belief state, Future Plan Predictor identifies most likely plans from the current belief state and constructs a tree of action-nodes, known here as a *plan-tree*, representing a set of planning paths highly likely to be taken by the user. An action-node includes a *query* for the information that is required for the action (*e.g.*, a database query), the *priority* (*i.e.*, the probability that the user will take the action), and a set of *constraints* (*e.g.*, a deadline constraint specifying the time by which the data must be retrieved). This plan-tree is then supplied to the Information Gatherer.

3.3 Information Gatherer

Given a plan-tree of predicted information-gathering tasks, Information Gatherer determines (or schedules) when and which information sources to use in order to satisfy the information needs of the user as well as coping with resource constraints (*e.g.*, network bandwidth) imposed by the problem domain; that is, the agent should not interfere with the user's planning activities by overusing computing resources. Initially, the information-gathering tasks are ordered by the priorities and the deadlines, ensuring not only the acquisition of the most useful information, but also a timely acquisition of data. In order to accommodate changing information requirements, Information Gatherer must optimize its current schedule incrementally to satisfy newer (thus more relevant) information-gathering constraints. The retrieved data is stored locally until used by Information Presenter.

3.4 Information Presenter

The agent directly interacts with the user through Information Presenter, which selects a subset of data from the locally cached data, and presents to the user at appropriate times. When to present which data is determined by the estimated user's future information needs. In order to avoid information overload, Information Presenter must only present data in temporal proximity to the actual need, with a sufficient time for the information to be useful for the action at hand. Additionally, Information Presenter must select an appropriate presentation format when offering information to the user. Finally, user feedback (*e.g.*, whether the presented information has been used) is collected and is provided for the agent as reinforcement in order to allow future improvements on the quality of supplied data.

4 Related Work

Plan recognition has been studied in various fields: assistive technologies where assistant agents can guide a human user to execute a plan correctly [6]; cooperative multiagent problems where individual agents can infer the plans of other agents to synchronize their actions [13]; adversarial multiagent systems where an agent tries to figure out the intention of an adversary from observed actions [3]; intelligent user interface that can predict the next user action [8], and more can be found in a survey as in [2].

There has been a renewed interest in decision-theoretic approaches to plan recognition, shown in recent efforts in which planning algorithms are applied to recognize user plans without the need for elaborate plan libraries of possible plan alternatives [12]. Notably, researchers in cognitive science use an MDP model similar to ours to represent how a human predictor recognizes the plan of another actor by observing a sequence of the actor's current activities [4].

In an approach known as imitation learning (*a.k.a.* inverse optimal control), an expert demonstrates to an (apprentice) agent a set of (semi-) optimal decision-making examples during a training session, from which the agent tries to learn the (hidden) reward function that matches with the expert's decision making [10, 1, 14]. This approach is generally restricted to cases where a reward from a state linearly depends on a set of features of the state. Our current implementation does not have the learning capability for updating the MDP user model online; the learning capability will be sought in future work.

A POMDP-based approach was used in [6] to assist dementia patients, where the agent learns an optimal policy to take a single best assistive action in a belief state. In contrast, the *ANTIPA* architecture separates plan recognition from the agent's action selection (*e.g.*, gathering or presenting information), which allows the agent to asynchronously plan and execute multiple alternative information-gathering (or information-presenting) actions. Our work is thus focused on how the agent can proactively assist the user after recognizing the user goals, *i.e.*, the agent can predict the user's most likely planning alternatives; subsequently, the agent can autonomously plan out assistive actions such as collecting relevant information.

Regarding information gathering, an approach exists for speculative plan execution in I/O-bounded planning problem domains [5]. In this approach, an agent uses a query classifier to speculate an outcome of a slow (computationally demanding) plan operator; continues to execute the next part of the plan using the guessed data; and verifies the speculated part of execution when the actual data from the source becomes available. In contrast, our information agent utilizes plan recognition techniques to predict the user's *future* plan and prefetches the information that the user is likely to need in the future.

5 Applications

The ANTIPA architecture is generic enough that, given the proper domain modeling, it can be used in many applications where information assistance in dynamic environments under time pressure is necessary. In this section we propose two conceptual applications of the ANTIPA architecture in the context of coalition operations. The first application is intended to be used for pre-mission planning, where the commander has to put together a plan of action while aggregating information relevant to the plan being developed. The second application is intended to be used by troop leaders on the ground to both track mission progress but also to proactively gather information needed during the mission.

5.1 Plan creation support

In the context of plan creation, we envision the applicability of the ANTIPA architecture in supporting a commander's planning process by prefetching relevant information during the planning process and presenting it to the commander. An example of existing system in which ANTIPA could be integrated is TIGR [9]. TIGR integrates a map interface with user-created content collected during operations in an area. It helps troops disseminate fine-grained intelligence collected in the field and pass this experience on as troops are rotated through an area of operations. Nevertheless, one of the potential side-effects of unlimited user-created content is the need for some kind of relevance filter so that a user does not need to sift through vast amounts of potentially unrelated information, presenting only what is needed for a particular part of the plan creation process.

Besides helping a commander cope with information overload from a single source, ANTIPA can also manage and query multiple information sources that may be relevant in mission planning. This is of particular use in coalition operations whereby multiple forces must interact and be made aware of each other's movements and plans. Here, ANTIPA can be used to communicate with other software agents being used by different forces in order to facilitate team operations.

5.2 Plan execution support

During plan execution, the predictive capabilities of the ANTIPA architecture can be used during a mission to detect the need for and proactively start collecting information that only becomes available after a plan of action started being carried out, or information that needs to be collected in close temporal proximity to its usage. For example, ground troops about to move through potentially hostile territory might need up-to-the-minute aerial imagery (*e.g.*, from unmanned aerial vehicles) over the movement area. However, in order for this imagery to be relevant in detecting potential threats, it must be collected and used immediately before troops enter the area. Moreover, the quality and utility of aerial imagery depends on the positioning of the collecting aircraft, so there is necessarily a time lag between the decision to collect imagery and the maneuvers necessary to accomplish this.

ANTIPA can also be used to bring selected TIGR information at the most relevant moments during a mission to a PDA-like device. In this scenario the agent can predict possible destinations for troops using some kind of mission profile, and proactively consult and present relevant information to the user without the need for a user to create complex queries using the PDA interface. A possible mock up interface for this application is illustrated in Figure 2.

6 Conclusions

In this paper we have reviewed the *ANTicipatory Information and Planning Agent* architecture [11] and proposed a vision for applications of the ANTIPA architecture in coalition operations. These applications range from pre-mission planning support to up-to-the-minute plan execution support. Research on ANTIPA has been conducted in the context of the UK Ministry of Defence

Figure 2: Mockup user interface for a ground commander application.

(MoD) and US Army Research Lab (ARL) *International Technology Alliance*[2].

We have developed a fully functional prototype using a decision-theoretic approach that has been tested in human experiments using a simple planning game yielding promising initial results. Our current efforts aim at refining the information-managing components. Future work will include enhancing the information-gathering scheduler to take into consideration redundant information sources, as well as the design of a finer-grained process to reason about user cognitive overload. Moreover, we aim to develop user experimentation simulations based on realistic military scenarios.

Our work is novel in utilizing plan recognition techniques in speculative plan execution where the agent prefetches information that a user is likely to need in the future. The notion of speculative plan execution was introduced in [5] for I/O bounded problem domains where information needs are cross-dependent in a hierarchical manner (*e.g.*, a set of districts depends on a city to which they belong, a set of cities depends on a state, etc.); in this approach, information needs are predicted based on previous user queries. Instead, the *ANTIPA* agent predicts potential information needs by recognizing the user's future plans by observing the current activities.

Acknowledgement: Research was sponsored by the Army Research Laboratory and was accomplished under Cooperative Agreement Number W911NF-09-2-0053. The and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation here on.

[2] `http://usukita.org/`

References

[1] P. Abbeel and A.Y. Ng. Apprenticeship learning via inverse reinforcement learning. In *Proc. ICML*, 2004.

[2] M. G. Armentano and A. Amandi. Plan recognition for interface agents. *Artif. Intell. Rev.*, 28(2):131–162, 2007.

[3] D. Avrahami-Zilberbrand and G. A. Kaminka. Incorporating observer biases in keyhole plan recognition (efficiently!). In *Proc. AAAI*, 2007.

[4] C.L. Baker, R. Saxe, and J.B. Tenenbaum. Action understanding as inverse planning. *Cognition*, 31:329–349, 2009.

[5] G. Barish and C. A. Knoblock. Speculative plan execution for information gathering. *Artif. Intell.*, 172(4-5):413–453, 2008.

[6] J. Boger, P. Poupart, J. Hoey, C. Boutilier, G. Fernie, and A. Mihailidis. A decision-theoretic approach to task assistance for persons with dementia. In *Proc. IJCAI*, pages 1293–1299, 2005.

[7] A. Fern, S. Natarajan, K. Judah, and P. Tadepalli. A decision-theoretic model of assistance. In *Proc. of AAAI*, 2007.

[8] P. Gorniak and D. Poole. Predicting future user actions by observing unmodified applications. In *Proc. AAAI*, pages 217–222, 2000.

[9] Mari Maeda. Tactical Ground Reporting (TIGR) System — Advanced Soldier Sensor Information System and Technology (ASSIST). Online, available at `http://www.darpa.mil/ipto/programs/assist/assist_tigr.asp`, 2010.

[10] A.Y. Ng and S. Russell. Algorithms for inverse reinforcement learning. In *Proc. ICML*, pages 663–670, 2000.

[11] Jean Oh, Felipe Meneguzzi, and Katia P. Sycara. ANTIPA: an architecture for intelligent information assistance. In *Proceedings of the 19th European Conference on Artificial Intelligence*, page (to appear), 2010.

[12] M. Ramírez and H. Geffner. Plan recognition as planning. In *Proc. IJCAI*, pages 1778–1783. Morgan Kaufmann Publishers Inc., 2009.

[13] G. Sukthankar and K. P. Sycara. Robust recognition of physical team behaviors using spatio-temporal models. In *Proc. AAMAS*, pages 638–645, 2006.

[14] B.D. Ziebart, A. Maas, J.A. Bagnell, and A.K. Dey. Maximum entropy inverse reinforcement learning. In *Proc. AAAI*, pages 1433–1438, 2008.