Perspectives on the State of Modeling and Simulating Human Intent Inferencing

Sheila B. Banks, Ph.D.
Air Force Research Laboratory
Air Force Agency for Modeling and Simulation
12350 Research Parkway
Orlando, FL 32826
sbanks@calculated-insight.com

Abstract

This paper presents the current state of research and practice for human intent inferencing, recognition, and application. The paper addresses both the specific topic of human intent inferencing and areas closely associated with human intent inferencing drawn from the research areas of artificial intelligence, human computer interaction, and cognitive science. Examples of these closely related topic areas are cognitive modeling and knowledge representation, plan/goal recognition and interpretation, and mixed initiative and collaborative problem-solving. Finally, the paper closes with a summary of the issues still remaining in the area of human intent inferencing and some recommendations on future research directions.

Introduction

As work and tasks become more complex and users come to require and rely ever more upon a larger number of data resources, their current means for accessing the data and performing the tasks are increasingly becoming perceived as being inadequate, wasteful, and even counterproductive. Clearly, the traditional cycle for accomplishing tasks that require computer support generally impedes efficient completion of a task. This cycle generally requires that the user develop a mental model for accomplishing the task, then determine the sources to be accessed, search the sources sequentially, analyze and correlate the results with some aid from the computer, and then present the results. This cycle is also increasingly viewed as being too sequential and vulnerable to disruption to permit the effective and rapid accomplishment of work in today's complex and dynamic world. When we couple the realization of the poor quality of today's work approach with the necessity for concurrent execution of tasks, frequent disruptions of task execution, tasks of varying duration and complexity, and the increasing demand for novelty in task solution, the argument for user assistance in task performance becomes compelling. To provide effective assistance, a computational agent must use a model of the work and work tasks and of the human user along with the ability to recognize user work progress and task accomplishment, current user focus, and projected user activity within the system and use this recognition to act in anticipation of the user's needs (Banks and Lizza, 1991). In other words, the ability to infer the user's intent and use this user intent determination to drive system resources is required in today's complex environments. It is this last topic, user intent and intent inferencing and how we can continue to draw from other fields and technologies to improve it (Llinas and Deutsch, 1999) that is the focus of this paper.

The motivation for this paper comes from the inadequate state of current research and its capability for advancing the state of current technology for intent inferencing. User intent inferencing can be used to support execution of user actions and work execution including the ability to assist the user in his/her thought processes. However, the main issue of intent inferencing, the recognition of the user's plan and conclusion of future user activity by some computational method, remains because we lack a proven methodology to bridge the gap between plan recognition and determining the user's anticipated problem solution path. Currently, the computational paradigm favorite is intelligent agent technology where an intelligent agent is a computer entity that collaborates with and helps a user. The roles of an intelligent agent include perception of dynamic conditions in the environment; action to affect the environment; and reasoning to interpret perceptions, solve problems, draw inferences, and determine actions. The computational method (intelligent agent) must determine the user's solution to plan and work accomplishment in spite of the fact that the user does not explicitly communicate the plan to the computer. Instead, the intelligent agent must rely upon the "conversation" between the computer and the user to provide the agent with the information it needs to determine the user's intent and thereby to anticipate user needs in relation to the work tasks that the user is performing. The interaction between human and computer also allows the agent to assemble a model of the user's behavior. Currently, intent is best inferred within the context of highly repetitive tasks and poorest for high-level, novel cognitive tasks, which are the very ones that require novel approaches or ideas. However, the potential for intent inferencing lies far beyond the current feeble state and can play a key role in employing technology to improve human performance because it can enable better decisions, better modeling and understanding of user activity, and overall increased efficiency for the user. Intent inferencing has the potential to be used to minimize disruptive breaks in the user's thought processes and to minimize the number and magnitude of cognitive disruptions that a user must deal with when accomplishing work tasks. Intent inferencing can also serve as a short-term memory aid to maintain the status of each work task that the user is engaged in (and so minimize the severity of the disruption in work execution) and also to highlight action options to the user as well as to assist the user in the performance of tasks. However, these benefits will not be realized without additional research to lay and extend the foundation for intent inferencing representation and implementations; research that must touch upon the fields of artificial intelligence (AI), user modeling, human-computer interaction (HCI), and cognitive science (Llinas and Deutsch, 1999). The research that is called for will not only have to address these technologies that support intent inferencing but also techniques for melding the technologies effectively. Today, we lack the requisite technological insight to achieve this capability.

The remainder of this paper is organized into two main sections. The next section contains a brief overview of intent inferencing and a brief look at a few current intent inferencing research projects. Section Three contains a discussion of the research directions and options that should form the basis for future intent inferencing research. Following these two main sections, the paper concludes with a short summary.

The Current State of Intent Inferencing

As a prelude to making recommendations for further research in the area of intent inferencing, this section briefly examines the current state of intent inferencing from the point of view of relevant technologies and research projects. The section opens with a review of relevant technologies, then continues on to a discussion of some intent inferencing research projects, and closes with a short assessment of the shortfalls in our knowledge about intent inferencing.

What is Intent?

Since the focus of this paper is on the state of the art for intent inferencing, it is appropriate to specify exactly what is meant by user intent. Ntuen defines human (user) intent as "mental states which drive actions" (Ntuen and dePontbriand, 1997). Geddes (Geddes, 1994) defines intent determination as the process of forming a causal explanation of the observed modes of action in terms of the observed modes of action. Jones advances a different definition of intent in his paper (Jones, 1988). Jones defines intent determination as the process of understanding operator actions in order to infer the operator's intention in executing those actions. Cohen (Cohen, 1990) takes a different approach to defining intent; in his paper, he argues that intent guides future planning and constrains other intentions, with the content of intention being a self-referential representation for its own conditions for satisfaction; with intent being tied to commitment to a current action which serves, in turn, to limit the spectrum of future choices (intentions) available. One final example intent inferencing definition is an approach used by Brown whose approach to predicting user intent identifies the salient characteristics of the domain environment and specifically determines the goals a user is trying to achieve (Brown, Santos, and Banks, 1998). Brown's approach is based on an intent definition that states what a user intends to do in an environment is the result of events occurring in the environment and the goals he/she is trying to obtain as a reaction to stimuli. These goals can be explicit (e.g., opening a file for editing) or implicit (e.g., reducing the workload). To achieve a goal, a user must perform certain actions to achieve the goal. The need to model and understand user intent arises in a wide number of research and development venues and, as a result, the user intent research community has yet to settle upon a single definition for user intent. However, given the breadth of applications of user intent, this ultimate, single definition for user intent may not be possible.

Background for Intent Inferencing: User Modeling

One research field integrally related to intent inferencing is user modeling (Banks and Stytz, 1999a; Banks and Stytz, 1999b; Banks and Stytz, 2000; Banks and Stytz, 2001; Jones, Chu, and Mitchell, 1995; NATO RTO-047, 2001; Riley, 1992; Rouse and Morris, 1996; Sheridan, 2000). The field of user modeling can be considered as a marriage between artificial intelligence (AI) and human-computer interaction (HCI). User modeling is traditionally concerned with how to represent user knowledge and interaction within a system with the purpose of adapting that system to the needs of the user and with assisting the user. Within user modeling, user intent can be defined as mental states that drive actions or actions a user intends to perform in pursuit of a goal state. User modeling and the run-time determination of user intent are exercised within both AI and HCI to drive system information processing, decision support processing, and interface requirements processing.

Many types of user models have been developed for a variety of purposes. Each type of model represents specific attributes of the user of a computer system, and each is useful in applications for which it was designed. However, no model represents everything and investigators with one set of aims may find models useless that were devised for other purposes. For user models, there is a distinction between competence models (those which determine what a user could do) and performance models (those which determine what a user is likely to do). Physical models, based on empirical knowledge of the human motor system and focus on task execution, are examples of competence models. The two user models of primary interest, behavioral and cognitive models, are both performance models in that they are used to determine a user's future actions. The primary difference between the two models lies in the level to which the user is modeled. Both models observe the user's execution of actions; however, cognitive models then attempt to determine the user's goals, whereas the behavioral model directly forecasts user activity. User modeling research may also factor addition aspects into the model such as demographic factors (age, gender), professional factors (expertise level), physiological factors (reaction, workability), and psychological factors (understanding, memory, working memory, or cognitive load) (Banks and Lizza, 1991; Banks and Stytz, 1999a; Banks and Stytz, 1999b; Jones, Chu, and Mitchell, 1995; Riley, 1992; Rouse and Morris, 1996; Sheridan, 2000).

Human cognitive models have been studied by researchers in the field of psychology for many years (Eggleston, Young, and McCreight, 2000). Cognitive psychology is concerned with understanding tasks in which a stimulus is processed in some way before a response is chosen. The underlying assumption of much of cognitive psychology is that a human perceives the world and produces some representation of it in his or her mind, called the "problem space" or "mental model". This representation is what we would usually call "knowledge." This knowledge may be described in terms of the concepts that we possess, the relationships between those concepts, and our capacity to make use of those concepts. The human then manipulates the knowledge and produces some output, or behavior, that can be observed. Humans form cognitive models of their environment to make sense of and organize the information they observe. Similarly, a computer system may also use a cognitive model of its environment and its user as it determines how to assist the user. Cognitive models represent aspects of users' understanding, knowledge, intentions and processing, and tend to have a computational flavor. Put simply, a cognitive model represents the human user as a collection of goals and a set of actions to accomplish the goals.

A human behavioral model, another performance type of user model, is intended to address some of the difficulties encountered when using a cognitive user model. In a behavioral model, the behavior of a system is manifested in input-output relationships; the user's behavior can be defined as a succession of states. Put another way, a behavioral model represents the human user as a collection of sequences of actions that the user performs. This model observes and predicts the activities of the user. The system does not attempt to determine the user's goal, as done with a cognitive model, but directly predicts future user actions based on the status of the environment and past user actions.

Background for Intent Inferencing: Ontologies

Another integral technological foundation for intent inferencing is that of ontologies. One of the keys to effective intent inferencing and to laying the foundation for developing the other technologies needed for intent inferencing is the development of ontologies for users' work domains (Chandrasekaran, Josephson, and Benjamins, 1999; Everett et al., 2002; Gruber, 1993; Guarino and Welty, 2002; Holsapple and Joshi, 2002; Kim, 2002). There are a number of different but mutually supportive definitions for an ontology; an ontology can be defined as a formal explicit specification of a shared conceptualization; or as an explicit specification of an abstract, simplified view of a domain to be represented; or as a specification of the concepts in a view and relationships in a domain.

An ontology consists of terms, their definitions, and axioms relating to them. The terms are normally organized in a taxonomy. An ontology allows intent inferencing to span threads of work and can minimize the opportunity for mis-communication and misunderstanding between agents and between the agent system and the human user. An ontology is a crucial foundation for communication between the representation of the user to the computer system and the representation of the system work to be accomplished given user direction. The ontology allows a computational agent to understand user work processes and then allows the computational agent to translate this work process user input into the computer system to direct system resources and processing.

Intent Projects

A few projects have been undertaken with the objective of developing methodologies and technologies for developing intent models and for incorporating user intent modeling into the functionality of agent systems and other types of decision support systems. In this subsection, we will briefly review some of these efforts.

One approach to modeling user intent is the use of plan-goal graphs as described by Geddes (Geddes, 1994). The use of plan-goal graphs that Geddes describes relies upon a directed acyclic graph that is an abstraction class hierarchy and a compositional class hierarchy that represented the purposes of the system. Geddes claims that a goal represents a specific criterion on the state of the system that can be tested to determine if the criterion has been satisfied by examining the state of the system. The most abstract concepts in the plan-goal graph are the top-level nodes. Plans represent actions that can be performed to achieve goals. The most tangible or concrete goals are at the bottom of the graph and represent specific actions. In Geddes' approach, intent interpretation is accomplished by trying to find a path through the plan-goal graph from an observed instance of an action to a plan or goal that was inferred previously. The basic paradigm is the linking of the observations to modeled goals in the system. Geddes, in a follow-on paper (Geddes and Lizza, 2001), presents a discussion of the types of services and capabilities that can be provided by a system that employs intent as a component of user support as well as an overview of applications that they have developed that have intent inferencing as a component.

Another project that has employed intent inferencing as a means for supporting user action is the work conducted by Callantine (Callantine, 2001; Callantine, 1999), which describes methods and applications for intent inferencing for air traffic controllers using a Bayesian network model and a Bayesian network temporal model. They employed a software framework of their own design to develop intent-aware intelligent agents that support air traffic controller teams and employ a model of perception for the air traffic display.

Lesh (Lesh, Rich, and Sidner, 2001) presents an overview of research that has been conducted at their laboratory that is related to intent inferencing in support of one-on-one collaboration and also describes projects that have been developed that use the intent inferencing capabilities. Another project that employs intent inferencing to aid users is described by Fano (Fano, 2001); the objective of this project is to improve the poor services typically provided by mobile devices by enabling collaboration between a user's device and the electronic services surround.

Kambhampati (Kambhampati, Mali, and Srivasiava, 1998) describes how they extended their previous work on refinement planning to include hierarchical planning by providing a generalized plan-space refinement that is capable of handling non-primitive actions. The generalization provides a means for handling partially hierarchical domains, while accounting for the user intent that is inherent in the activity.

Gonzalez (Gonzalez and Sacki, 2001) describes an approach for representing tactical decision-making in what they call autonomous intelligent platforms using their competing context concept. A context is used to control the operation of a platform, and there is always a context in control, allowing the platform to be controlled via transitions between various contexts. Their competing context concept introduces a way to control the transition process without the need to predetermine the next context to be performed.

Franke (Franke, Brown, Bell, and Mendenhall, 2000) describes their work to help users by inferring intent when performing tasks. They contend that applications that have access to user intent and task context can support better, faster decision-making on the part of the user. In support of this contention, they developed AUTOS, an approach to the implementation of individual and team intent inference. AUTOS employs observable contextual clues to infer current operator task state and predict future task state. Using the concepts of activity theory, the authors contend that AUTOS task models can be hierarchically organized to infer team intent.

Shortfalls in Intent Inferencing Capabilities

Upon review of the current intent inferencing research and application efforts, a few salient features concerning the state of intent inferencing research become apparent. Obviously, and unfortunately, there are and have been relatively few research projects undertaken and, by and large, these efforts have been secondary efforts within larger research projects and an overall coordinated strategy to deal with the intent inferencing research requirements is not evident. However, in spite of the relative lack of research in intent inferencing, the level of desired capability for intent inferencing systems remains high but real world user expectations cannot be met. As a result, it is clear that there is a large gap between desired capabilities and current technology coupled with an overall ineffective approach to remedying the shortfall. Until these shortfalls are addressed in a systematic manner within a coordinated suite of research thrusts, systems that rely upon a determination of user intent in order to perform adequately will continue to fall short of user expectations and requirements.

Future work tasks will suffer these same shortfalls even considering the anticipated technological advances. The increasing amounts of information to be accessed, correlated, and processed coupled with an increasing number of work tasks to be performed indicates that the demand for computational agents using intent inferencing to operate in anticipation of user needs and provide user assistance will grow. However, unless we begin to develop an understanding of how to infer user intent and employ it to help a user accomplish real work, the increased computing power will not be effectively exploited because the user will still be required to deal with the cognitive burdens of task context changes and unassisted location and integration of information. In short, increased computing power will not result in improved user performance unless some of that power is effectively used to ease the cognitive demands placed upon the user during the course of performing work.

Moving Toward the Future: Recommendations

Before turning to specific recommendations, we should note that at the heart of intent and intent inferencing are agent systems and user models, a computable understanding of the activities to be performed, and the allocation of activities between the user and the intelligent agent system (Burstein and McDermott, 1996; Burstein, Ferguson, and Allen, 2000; Llinas and Deutsch, 1999; Norman, 1997; Schneiderman, 2002; Sheridan, 2000). The allocation of activities between the computer and human is especially difficult because user needs for support are complex, varied, sometimes overlapping, redundant, concurrent, and almost always non-sequential; therefore, function allocation must address these concerns (Geddes, 1997). As Sheridan points out (Sheridan, 2000), we are still unable to determine how to perform function allocation because of differences in understanding of user's thinking, interactivity, and purpose as well as a number of trends including the following: continued development of technologies that promote human supervisory control of a system, the insight that the proper allocation of activity changes as the task proceeds, that uncertainties remain concerning normative behavior for user performance of tasks, the insight that function allocation is not just a run-time issue but extends back to system design, and the continuing difficulty in determining the technology that determines a useful function that is compatible with human psychological orientations toward machines.

In support of intent inferencing, an intelligent agent system can be employed because of the complexity of the tasks that require assistance or that could be allocated to the intelligent agent system. An agent system also naturally supports the user's work domains and permit seamless performance of concurrent tasks with the agent system operating in anticipation of user needs and work focus changes, in spite of disruptions, so that the user's cognitive load is minimized and so that the user can maximize attention on the task. The intelligent agent system must know how to find and retrieve data that is relevant to the user, analyze the data, determine intent, and present relevant data to the user in a manner that aids the user in performing the task(s) at hand. The chief technical focus and challenge in this cycle is determining user intent. Intent inferencing is used by the intelligent agents to direct the system to operate in anticipation of user needs and to permit coordination of effort within and across work tasks. To perform at this level of support, the agent system must determine user goals, plans, and focus of attention. In addition to the technical challenges posed by determining user intent, user privacy and security issues must be addressed and steps must be taken to insure that information related to user tasks and data sources does not propagate through the system; however, addressing these concerns lies beyond the scope of this paper.

The gaps in our capabilities for intent inferencing indicate where we should begin and how we should organize our research thrusts to enable improved determination of user intent. At a basic level, it is clear that we do not yet have a comprehensive definition or clear understanding of intent inferencing or a baseline theory for intent and intent inferencing. Intent can be viewed as future intention embedded within current intention and activity, or, alternatively, as the result or outcome of a set of beliefs and desires on the part of a user. However, it can also be viewed as a conscious process, a state of the nervous system, or even patterns of behavior. One challenge that lies before us, then, is to unify the field and come to a common understanding of what we mean by intent and its processes. To frame the following discussion, it has become apparent that to effectively enable intent inferencing, a system must perform two tasks well: 1) the agent system must be able to acquire information from the user and 2) the system must be able to analyze and manage the information it acquires. In light of these two basic performance requirements, the system must contend with the missing data problem, which arises because the user does not explicitly communicate goals, plans, tasks, and focus of attention to the computer/agent system. As a result, the computational agent system must make estimates to fill in the gaps in the information and so support the user.

One organizing framework that should be used to permit the agent system to understand and support the user's work is the concept of the work thread. A work thread is a connected series of work activities directed toward the attainment of a goal. Each activity in a thread can consist of multiple steps (or actions) with some steps being purely human thought and others being activities and data access actions. This conceptual approach to problem space decomposition along with other approaches like aspect oriented programming (Elrad, Filman, and Bader, 2001) may prove useful in allowing us to make advances in portions of the intent inferencing domain and also permit us to frame experiments that are constrained enough to provide objective results and measures of performance.

One aspect of intent inferencing that calls for extensive research is the topic of support for the user when the user addresses a new task. Research is needed in order to develop the capability for an agent system to infer intent in a user task environment where ad hoc solutions are developed, where there is no prior experience in the exact task, and where non-linear thinking and intuitive exploration of the problem space and solution space are prevalent. Intent inferencing will be particularly difficult to provide in these situations since the organization of the task is uncertain and the necessary information and its organization is not known. The support needed appears to be required along the lines of assisting the user with brainstorming ideas and in the presentation/evaluation of creative options. The model representations, reasoning processes, and knowledge needed to enable intent inferencing that facilitates creative user activity are broadly open questions and research is required along several dimensions.

In a related vein, intent inferencing is necessary for collaboration among team and group members as collaboration is increasing in importance for users in task accomplishment within many work domains. In this aspect of work, intent inferencing is needed in order to identify the people who may be useful to the user in accomplishing a task and also to identify other agents or agent systems whose support/collaboration may be of use in accomplishing the task. In addition to support novel tasks, research in intent inferencing that supports common collaboration task execution is required. Our current technology state provides little advancement for intent inferencing for collaboration work even for common or repetitive tasks; therefore, basic research to lay the foundation for effective intent inferencing to identify and initiate collaborative opportunities for common task performance is needed.

To advance the current state of intent inferencing and develop improved intent inferencing capabilities requires investigation and use of techniques and technologies for user modeling and human behavior representation. In other words, to model human behavior is to establish a framework in which to perform intent inferencing and improvements in user modeling provide enhanced capabilities for deriving user intent. To improve user modeling, for example, current rule-based schemes must incorporate additional techniques to enhance robustness, like uncertainty reasoning incorporating fuzzy logic or Bayesian techniques. In addition, to satisfy future requirements for user modeling, a common framework for understanding and development is needed (Banks and Stytz, 2001; Banks and Stytz, 2000). The development environment should be a component-based framework providing the opportunity to create, develop, test, and validate user modeling techniques and the models themselves. The development environment should unify the research, development, and transition efforts of user modeling research and focus on the requirements for human behavior modeling within prototype domains. The development environment should facilitate achieving two objectives: 1) the integration of existing models and modeling techniques into the environment to allow for user assisted model integration, testing, and analysis; and 2) the development of new models and modeling techniques to allow for user assisted, component-based model building to enable the testing of new methodologies/theories for human behavior modeling and to allow the seamless integration of the model development process with the model execution and performance within a selected environment.

Given the many approaches to knowledge representation and manipulation and the architectures for these knowledge processing systems, one can readily identify advantages with each scheme. On the other hand, each technique brings associated issues that must be dealt with in order to pursue the technology appropriately with hope of its use for user modeling and the representation of user intent in the future. One approach to deal with the shortcomings of various representations and their manipulation techniques and gain full advantage of the benefits of each is to implement the various knowledge representation schemes and their processing methods in a cooperative problem-solving architecture; namely, a hybrid architecture. A hybrid architecture would enable a variety of knowledge representation and manipulation techniques and allow the exploitation of a number of technologies. This type of architecture must, necessarily, also separate knowledge from processing and allow component-based development for both knowledge and processing opportunities. A hybrid approach enables more than one approach to be used at the same abstraction level and multiple abstraction levels. However, even with hybrid approaches, not every situation can or should be modeled in every type of knowledge representation and care must still be used to select the correct modeling technique appropriate to the modeling need. For example, within the hybrid architectural framework, neural networks and case-based reasoning could be applied appropriately to address the learning issue and additional research forthcoming could then be utilized for learning and memory to improve the operation of the agent system in inferring user intent. In addition, a hybrid approach to situation assessment that incorporates belief networks and fuzzy logic into rule-based processing could be advanced for the same purposes. The basis for development of a hybrid architecture for future user modeling and intent inferencing needs will probably necessitate integration of concepts from a number of current approaches with the extension of these concepts into the next generation cooperative modeling paradigm (Banks and Stytz, 2001).

To support the need to improve intent inferencing primarily within an intelligent agent system, ontologies that support the users' domains of operation are also needed. Within an intelligent agent system, an ontology is crucial to allowing the agent to organize and understand the activity and to support the user effectively (Put Reference here). Ontologies are needed for each domain of operation and also across domains of operation. The ontologies must be constructed so that they provide an agent system with the understanding needed to comprehend a domain of operation, to measure their own effectiveness of support for the user's operations, and for the agents to have a broad understanding of each user's domains of work. The ontologies must be broad enough to support the user as they cross domains and also contain enough content to allow an agent system to understand the nuance of work domains so that they can support creative and novel work. Finally, and perhaps most important for the broad application of intent inferencing for collaborative work, is that the ontologies must be constructed so that they can be expanded and refined by the agent systems as the agents learn about the user's work domains and the user's approaches to performing ongoing and novel work. As regards a user modeling common development environment, an ontology is foundational because it describes the domain and its important parameters in general grammar terminology and forms the basis for both user model integration and model interoperability so that all models integrated or developed in the environment may communicate on both the syntactic and semantic levels. The development of and adherence to an ontology within a common development environment will enable legacy models to be integrated by using the ontology to form the basis for model application program interfaces (APIs), which will enable new models to be developed within the environment to be interoperable with the integrated models.

One technology that could be used within the ontology effort to address the shortfalls in intent inferencing and to assist an agent in inferring intent is to use hypertext markup language (HTML) and eXtensible Markup Language (XML) (Copenkus and Hoodbhoy, 1999) tags to embed semantics concerning tasks and the data within the data being accessed and managed. Embedding semantics would allow the agent to not only know what the user is doing via monitoring of inputs by the user but would also give the agent some data concerning the meaning and definition of the information accessed by the user and potentially additional insight into the task(s) that the user is executing. To successfully employ this approach, an ontology for each task and data set will be required and there will also be a need for the capability for correlating the ontologies and the data to enable recognition of similar data items across problem domains and work activities.

In addition to the research initiatives discussed above, there are additional areas that must be addressed in order to achieve a solid, reliable, and forward-looking capability for intent inferencing. One primary topic to be addressed is the need for intent metrics and measures of performance. Our current capability in this area is that we have no means for measuring or assessing the performance of an intent inferencing system apart from user subjective assessments of the performance of the system. The need is for a set of objective, external assessment measures that can be applied during research, development, and operation that permit continuous evaluation of the system. This capability is necessary not only for being able to develop an intent inferencing capability but also to allow the system to learn and to adapt to each user and to continually improve its performance. In addition to a standard set of measures for the performance of an intent inferencing system, an independent, objective intent inferencing testing methodology is required coupled with a standard set of use cases (or evaluation scenarios) are needed. The testing methodology would employ the use cases to insure complete coverage and capability for an intent inferencing system, but equally important it would serve to support the assessment system for permitting evaluation of the capability of an intent inferencing approach and objective assessment of the performance of an intent inferencing system. Our point here is that the field of intent inferencing requires the capability for performing objective assessments of the performance of an intent inferencing system, and that this capability should have at its foundation a standard set of metrics for measurement of performance, a set of standard use cases that can be employed in conjunction with the metrics to assess the system in a standard, repeatable manner, and finally that there should be a test and evaluation methodology that uses the metrics and use cases to evaluate the intent inferencing system. Because of its broad use and associated standardization efforts, the Unified Modeling Language (UML) (Booch, Rumbaugh, and Jacobson, 1999) may be the best foundation for the documentation of use cases. Undoubtedly, the use cases will have to be application domain specific but there will be opportunities for reuse of scenarios and use cases, hence we believe that UML or a similar technology should be used for documentation in order to minimize the cost associated with the development of the use cases and also to permit the widest possible review of them. Given these testing and evaluation capabilities, the argument for the effectiveness and utility of intent inferencing can be made and supported and the field can make measurable progress in its ability to assist a user in the performance of tasks.

Summary and Conclusions

As a final motivation note, the ability for a system to determine user intent and utilize intent determination in further system processing is a crucial need for the development of any computational system that supports user assistance and collaboration in the execution of mixed initiative tasks in an adaptive, intelligent, learning work environment. The work system envisioned is symbiotic in that work tasks are broadly and interactively partitioned between the computer and the user, where the user, through an integrated natural language and graphical interface, provides guidance and insight into the broad classes of information that are needed to accomplish the task, and draws complex inferences from the data. The computer, on the other hand, performs computationally intensive tasks such as data acquisition, quantitative data analysis, and data mining/knowledge discovery to assist the user in accomplishing tasks. The system would use information visualization techniques to enable the user to comprehend the information and available options accessed and analyzed by the computer. The intelligent agents in the system would operate in anticipation of user information needs and ascertain user information requirements based on the current situation, utilizing a history of required information, and a user specific model to initiate data retrieval operations, present information, and assist the user in focusing on and analyzing the relevant information. The ability of the system to operate in anticipation of user needs and offer beneficial assistance hinges on a fundamental technical requirement: the representation of user intent and intent inferencing based upon user modeling techniques. This representation must incorporate an understanding of the user work objectives, goals the user is pursuing over time, past user actions, and a determination of what to improve over time to better model the user within the domain.

This paper has briefly reviewed the current state of research and practice for human intent inferencing, recognition, and application. The paper addressed both the specific topic of human intent inferencing and areas closely associated with human intent inferencing. Examples of these closely related topic areas are cognitive modeling and knowledge representation, plan/goal recognition and interpretation, and mixed initiative and collaborative problem-solving. The paper closed with a look at only a few of the issues still remaining in the area of human intent inferencing and some recommendations on future research directions. One question not yet addressed is the work areas and professions that may be most amenable to support and exploitation of intent inferencing. Areas such as information networks, aircraft flight, aircraft maintenance, and medicine may be the best areas to focus our research efforts since these areas have a wealth of commonly accepted terms, definitions, and practices, and as importantly, each of these professions strive to maintain clear and distinctive definitions for their words and tasks. Taken as a whole, a number of research efforts are required for intent inferencing to advance in its capabilities and for intelligent agent systems to provide effective and efficient support for user work support. This paper has pointed to a few of what should be the most important of these efforts.

References

Banks, S.B. & Stytz, M.R. 2001. "The Next Generation of Human Behavior Representation for Modeling and Simulation: The HBR Common Development Environment and Framework," *10th Conference on Computer Generated Forces and Behavioral Representation*, Norfolk, VA, 177-188.

Banks, S.B. & Stytz, M.R. 2000. "Advancing the State of Human Behavior Representation for Modeling and Simulation: Technologies and Techniques," *9th Conference on Computer Generated Forces and Behavioral Representation*, Orlando, FL, 87-100.

Banks, S.B. & Stytz, M.R. 1999a. "Inclusion of User Modeling Techniques for Human Behavior Representation in a Domain-Independent Computer Generated Force Architecture," *8th Conference on Computer Generated Forces and Behavioral Representation*, Orlando, FL, 457-468.

Banks, S.B. & Stytz, M.R. 1999b. "The Necessity and Development of User Modeling for Future Modeling and Simulation Systems," *21st Interservice/Industry Training Systems and Education Conference*, Orlando, FL, December, publication on CD-ROM.

Banks, S. B. & Lizza, C.S. 1991. "Pilot's Associate: A Cooperative, Knowledge-Based System Application," *IEEE Expert*, June, 18-29.

Booch, G.; Rumbaugh, J.; & Jacobson, I. 1999. *The Unified Modeling Language User Guide*, Addison Wesley, Reading, MA.

Brown, S.; Santos, E. Jr.; & Banks, S. 1998. "Utility Theory-Based User Models for Intelligent Interface Agents," *Lecture Notes in Artificial Intelligent 1418: Advances in Artificial Intelligent -- AI-98*, Springer-Verlag, 378-392.

Burstein, M.H. & McDermott, D. 1996. "Issues in the Development of Human-Computer Mixed-Initiative Planning Systems." In B. Gorayska & J. Mey (eds.) *Cognitive Technology*, Elsevier Science, 285-303.

Burstein, M.H.; Ferguson, G.; & Allen, J. 2000. "Integrating Agent-Based Mixed-Initiative Control with an existing Multi-agent Planning System," *Proceedings of the International Conference on Multi-Agent Systems.*

Callantine, T.J. 2001. "Air Traffic Controller Team Intent Inference," *AAAI Workshop on Intelligent Inferencing for Collaborative Tasks*, AAAI Technical Report FS-01-05, 32-39.

Callantine, T. 1999. "An Intelligent Aid for Flight Deck Procedure Performance: The Crew Activity Tracking System "Task Caddy," *Proceedings of the 1999 IEEE Conference on Systems, Man, and Cybernetics*, 1095-1100.

Chandrasekaran, B., Josephson, J.R., & Benjamins, V.R. 1999. "What Are Ontologies, and Why Do We Need Them?" *IEEE Intelligent Systems*, January/February, 20-26.

Cohen, P.R. and Levesque, H.J. 1990. "Intention is Choice with Commitment," *Artificial Intelligence*, 42, 213-261.

Copenkus, A. & Hoodbhoy, F. 1999. *Applied XML: A Toolkit for Programmers*. John Wiley & Sons: New York, NY.

Elrad, T.; Filman, R.E.; & Bader, A. 2001. "Introduction to Aspect Oriented Programming," *Communications of the ACM*, vol. 44, no. 10, October, 28-32.

Eggleston, R. G.; Young, M. J.; & McCreight, K. L. 2000. "Distributed Cognition: A New Type of Human Performance Model," *Proceedings of the 2000 AAAI Fall Symposium on Simulating Human Agents*, AAAI Technical Report FS-00-03, 8-14.

Everett, J.O.; Bobrow, D.G.; Stolle, R.; Crouch, R.; de Paiva, V.; Condoravdi, C.; van den Berg, M.; & Polanyi, L. 2002. "Making Ontologies Work for Resolving Redundancies Across Documents," *Communications of the ACM*, February, 45(2), 55-60.

Fano, A. 2001. "Mobile Valet: Enabling Collaboration between Remote Services, Mobile Users, and Task Locations," *AAAI Workshop on Intelligent Inferencing for Collaborative Tasks*, 98-105.

Franke, J.; Brown, S.; Bell, B.; Mendenhall, H. 2000. "Enhancing Teamwork Through Team-level Intent Inference," *Proceedings of the 2000 International Conference on Artificial Intelligence. IC-AI'2000, vol. 2*, Las Vegas, NV, 26-29 June, 1033-1039.

Geddes, N.D. and Lizza, C.S. 2001. "Practical Applications of a Real Time, Dynamic Model of Intentions," *AAAI Workshop on Intelligent Inferencing for Collaborative Tasks*, AAAI Technical Report FS-01-05, 40-45.

Geddes, N. 1997. "Associate Systems: A framework for human-machine cooperation," *Proceedings of the Seventh International Conference on Human-Computer Interaction*, San Francisco, CA, August, 237-242.

Geddes, N. 1994. "A Model for Intent Interpretation for Multiple Agents with Conflicts," *IEEE Conference on SMC*, San Antonio, TX, October.

Gonzalez, A.J.; Saeki, S. 2001. "Using Contexts Competition to Model Tactical Human Behavior in a Simulation," *Proceedings of Modeling and Using Context.: Third International and Interdisciplinary Conference, CONTEXT 2001*, Dundee, UK, 27-30 July, 453-456.

Gruber, T.R. 1993. "A Translation Approach to Portable Ontology Specifications," *Knowledge Acquisition*, 5(2), 199-220.

Guarino, N. & Welty, C. (2002) "Evaluating Ontological Decisions with OntoClean," *Communications of the ACM*, vol. 45, no. 2, February, 61-66.

Harper, K.A. and Zacharias, G.L. (2002) "Modeling Attention Allocation and Multitasking in Computational Human Behavior Representation," *Proceedings of the 11th Conference on Computer-Generated Forces and Behavior Representation*, Orlando, Fl, May 7-9, 59-66.

Holsapple, C.W. & Joshi. K.D. (2002) "A Collaborative Approach to Ontology Design," *Communications of the ACM*, February, 45(2), 42-47.

Jones, P.M. (1988) "Intent Inferencing by an Intelligent Operator's Associate; A Validation Study," *Proceedings of the Human Factors Society, 32nd Annual Meeting*.

Jones, P.M.; Chu, R.W.; & Mitchell, C.M. (1995) "A Methodology for Human-Machine Systems Research: Knowledge Engineering, Modeling, and Simulation," *IEEE Transactions on Systems, Man, and Cybernetics*, July, 25(7), 1025-1038.

Kambhampati, S.; Mali, A.; and Srivastava, B. (1998) "Hybrid Planning for Partially Hierarchical Domains," *Proceedings of the 1998 15th National Conference on Artificial Intelligence*, Madison, WI, July 26-30, 882-888.

Kim, H. (2002) "Predicting How Ontologies for the Semantic Web Will Evolve," *Communications of the ACM*, February, 45(2), 48-54.

Lesh, N.; Rich, C.; and Sidner, C.L. 2001. "Intent Inference Techniques and Applications at MERL," *AAAI Workshop on Intelligent Inferencing for Collaborative Tasks*, AAAI Technical Report FS-01-05, 46-57.

Llinas, J. & Deutsch, R. 1999. "Perspectives of Intent from Various Bodies of Literature and Perspectives for Developing New Models of Intention-Estimation and for the Development of Decision Aids," *Research Foundation of the State University of New York, Report CMIF-2-99*, Center for Multisource Information Fusion, February.

NATO RTO-TR-047, *Human Behavior Representation*. 2001. North Atlantic Treaty Organization, Research and Technology Organization, BP 25, 7 Rue Ancelle, F92201 Neuilly-Sur-Seine Cedex, France, July.

Norman, D.A. 1997. "How People Might Interact with Agents." In J. Bradshaw (ed.), *Software Agents*, MIT Press: Cambridge, MA.

Ntuen, C.A. & dePontbriand, Rene. 1997. "A Cognitive Network Simulation for Modeling the Commander's Intent," *The Institute of Human-Machine Studies*: Greensbore North Carolina.

Riley, V. 1992. "Human Factors Issues of Data Link: Application of a Systems Analysis," *Aerotech '92*, Anaheim, CA, 5-8 October, paper 922021.

Rouse, W.B. & Morris, N.M. 1996. "On Looking Into the Black Box: Prospects and Limits in the Search for Mental Models," *Psychological Bulletin*, 100(3), 249-363.

Schneiderman, B. 2002. "Understanding Human Activities and Relationships," *Interactions*, September/October, 9(5), 41-53.

Sheridan, T.B. 2000. "Function Allocation: Algorithm, Alchemy, or Apostasy?" *International Journal of Human-Computer Studies*, 52, 203-216.