

Phillips, J., Liebowitz, J. & Kisiel, K. (2001). Modeling the Intelligence Analysis Process for Intelligent User Agent Development, *Research and Practice in Human Resource Management*, 9(1), 59-73.

Modeling the Intelligence Analysis Process for Intelligent User Agent Development

Joshua Phillips, Jay Liebowitz & Kenneth Kisiel

Abstract

Intelligence analysis, whether competitive intelligence, business intelligence, criminal investigation, and the like, is a critical process in many domains. Unfortunately, there have been only limited models explaining the intelligence analysis process. We have developed a comprehensive, detailed model that steps the intelligence analyst through the process of eventually arriving at a conclusion. This paper explains this model which serves as a basis for encoding the model as an intelligent user agent to be used in conjunction with the Wisdom Builder knowledge management tool. This model should greatly help intelligence analysts in building knowledge management tools to help them in their investigative work.

Acknowledgements: We gratefully acknowledge the support of the Maryland Industrial Partnership System (MIPS) grant and WisdomBuilder LLC. for funding this effort.

Background

The purpose of this research project, funded by the Maryland Industrial Partnership System (MIPS) grant, is to build and deliver a client-based, Windows environment software application that acts as an intelligent agent to assist Wisdom Builder users in the criminal investigative domain. Wisdom Builder (2000) is a knowledge management tool that has been used in the intelligence and investigative domains to help users in the requirements, collection, analysis, and reporting phases.

Wisdom Builder supports activities across these four major phases of the analytical research process to help monitor areas of interest and develop strategies that promote innovation, productivity, and profitability. One of the main limitations of Wisdom Builder is that it may be somewhat difficult to use, partly due to its powerful features. In order to reduce the burden on the user and help guide the user through a session of Wisdom Builder, it would be helpful to have an intelligent user agent to provide recommendations to the user on performing his/her requirements and analysis functions using Wisdom Builder. In a sense, this intelligent user agent may be like a Microsoft Wizard to partly look over the shoulder of the user and provide suggestions on how best to gather requirements and perform the analysis steps. The agent resides in the background, monitoring the user's actions and offers suggestions and/or courses of action to the user based on the user's interaction with Wisdom Builder. At the user's request, the agent could interact with Wisdom Builder directly to perform the suggested actions. This research touches on personal agents (e.g., Gams (1996) and Soltysiak and Crabtree (1998)), software coaches, intelligent user interfaces (e.g., Intelligent User Interfaces Conference Proceedings (1999)), and case-based reasoning (e.g., Munoz-Avila, Hendler, and Aba (1999) and O'Leary and Selfridge (1999)).

The application domains selected for testing purposes are intelligence analysis and criminal investigation. In Heuer's book on Psychology of Intelligence Analysis (1999), the analyst typically applies six key steps in the analytical process: defining the problem, generating hypotheses, collecting information, evaluating hypotheses, selecting the most likely hypothesis, and the ongoing monitoring of new information. The FBI Academy's Special Agent's Training includes 16 weeks of intensive instruction in firearms, practical applications, physical training/defensive tactics, legal, forensic science, interviewing, informant development, communications, white-collar crime, drug investigations, ethics, organized crime, behavioral science, computer skills, and national security matters. A number of artificial intelligence-based systems like COPLINK (<http://ai.bpa.arizona.edu/coplink>), whose knowledge-based databases are used by the Tucson Police Department to provide large-scale intelligence analysis capabilities including the identification of previously unknown relationships, have been built to assist in the law enforcement area. However, a strong need exists to develop intelligent agent-assisted performance support systems to aid the investigative analyst in performing his/her critical functions.

The Wisdom Builder Intelligent User Agent

To help elicit the analyst's requirements for an intelligent agent for the Wisdom Builder tool, a web-based survey was used to ask the following major questions:

- name the top three features of Wisdom Builder that you find most helpful to you in your analysis;
- name three ways that could maximize the usability of these features;
- name some features that you would like to have in Wisdom Builder that currently aren't available to you.

From the results of this survey, the main features that the users wanted to have in the intelligent user agent were: having the agent look over the analyst's shoulder to make sure that he/she wasn't omitting useful facts and hypotheses in solving a case; making sure that the hypotheses and resulting conclusion seem consistent and reasonable; helping the analyst step through the analysis phase and key strokes in the Wisdom Builder product; having the agent help the analyst in the thinking and reasoning processes involved in making intelligence judgments.

After collecting these user requirements, the next step was to design the intelligent user agent. At first glance, it appeared that the Wizard technology may be very suitable for the agent. For example, in Microsoft Access, there are various Wizards such as the OutputWriter Wizard, ChartExcel Wizard, ReportExcel Wizard, NotesTable Wizard, Help Wizard, ReportRunWizard, PrintExcel Wizard, CodeBox Wizard, and Renaming Wizard. In Atkins (2000) article on generating Microsoft Wizard interface, a wizard is broken into two groups: (1) layout and (2) functional. The main layout features of a wizard are: Action buttons at the bottom of the window, Graphic in an area on the left, Data area to the right of the graphic and above the buttons, and Instructions displayed somewhere on the window. In addition to these layout features, a wizard interface often has the following functional features (Atkins, 2000):

- buttons are used to move through a set of input pages
- instructions change for every page
- each page is validated before continuing to the next page
- each subsequent page uses previous page's data
- nothing is formally committed until the whole process is complete, and the user can press the "cancel" button at any time and leave the wizard without leaving partial or invalid data
- the action buttons are activated/deactivated, or the label changes based on the page context
- the user presses the "finish" button.

In many of the Microsoft Wizards, case-based reasoning (CBR) is used. CBR involves analogical reasoning whereby a new situation is compared with existing cases in a case base, matching similar case features or adapting from those cases for resolving the new situation. Bayesian user modeling has also been applied and integrated with case-based reasoning systems to infer a user's needs by considering a user's background, actions, and queries. For example, the Lumiere Project by Microsoft Research applied Bayesian user modeling and served as the basis for the Office Assistant in Microsoft Office's suite of productivity applications (Horvitz et al., 1997). At the Navy Center for Applied Research in Artificial Intelligence at the Naval Research Laboratory, a system called INBANCA (Integrating Bayes Networks with Case-Based Reasoning for Planning) has been discussed (Aha and Chang, 1996).

Case-based reasoning was selected as the intelligent system methodology for the proposed agent. After trying Esteem's/SHAT's CBR Express and TecInno's CBR-Works (Germany) case-based reasoning tools, CBR-Works was chosen for the following reasons: (1) an executable/client version as well as a server version of the case-based application could be created, (2) the user interface handled some natural language processing, and (3) the tool had been used by a number of major companies, and (4) it was a fairly easy-to-use tool. The first part of this CBR application focused on answering user queries in setting up their application.

We first created a knowledge taxonomy that related to the main functions and tabs in Wisdom Builder (e.g., Table tabs, connection tabs, timeline tabs, link notebook tabs, create tabs, deliver tabs, etc.). Then, a listing of the typical user questions as related to each function/tab was created, and about 95 cases were then inputted into the case base which would respond to the answers associated with these questions.

The server version allows Wisdom Builder users to send additional questions and possible answers to Wisdom Builder Inc., who could then validate the information before confirming its acceptance into the case base. This server version also promotes an online community of practice of Wisdom Builder users worldwide.

This CBR approach helped us in better understanding the necessary processes and requirements, as related to Wisdom Builder user difficulties, and the general intelligence analysis methodology. This, in turn, facilitated the development of a comprehensive intelligence model which has been encoded into a tool called POINT (Problem Organization Intelligence Tool), developed by the authors in Visual Basic and MS-Access. The rest of this paper details this model.

The Competitive Intelligence Model

The next step was to further develop this agent by providing a capability to monitor the user's actions and provide suggestions during the competitive intelligence process. In order to provide this level of intelligence, the knowledge model of a typical intelligence operation had to be developed.

Both interviews and a literature review were used in order to create the model. Typical knowledge elicitation and modeling methodologies were only partially applicable since the problem domain was so abstract. However, an attempt was made to follow the general form of the "Task Analysis Worksheet" that is part of the CommonKADS methodology (Schreiber et al., 2000).

For each task that was identified within the larger analysis process, we explored the following aspects:

- The goal of the task
- The inputs and outputs
- The structures that we manipulated by the task
- Pre-conditions and post-conditions
- The agents involved
- Factors for judging successful completion of the task

The literature review was a more practical method for eliciting knowledge about such an abstract domain than interviewing. Because of this abstractness, it was exceedingly difficult to formulate interview questions that would elicit the above aspects and at the same time not be too broad or narrow. Our model is largely a synthesis of components that can be found in Friedman (1997), Barndt (1994), Kahaner (1996), Meyer (1987), and Heuer (1999).

The terminology and structures that we have adapted should provide ways to characterize an instance of an intelligence operation. If we can characterize aspects of an intelligence operation, then an intelligent agent that utilizes CBR might reason about past operations and offer useful direction to the user about the current operation. The discussion that follows describes what someone involved in CI (competitive intelligence) might be expected to be able to do. The focus of the model is on the process itself and makes no commitment to any computerized system.

General Phases in Intelligence Analysis

The process of intelligence analysis might be divided into an arbitrary number of steps. The Wisdom Builder software divides the process into four steps: Requirements, Collection, Analysis, and Report. This is probably the closest that one could want to come to a consensus in the field. However, the results of the survey and information gathered during our literature review suggest that intelligence providers might benefit from a more structured approach that would divide the process into more cognitively manageable chunks. Two of the main problems that intelligence providers have are: formulating the requirements that direct the operation; generating and evaluating multiple hypotheses.

Our model consists of seven phases that can be completed sequentially and iteratively. In addition, we identified four "mini-phases" that can be invoked during the main phases in much the same way a computer program invokes a function or method. The main phases are:

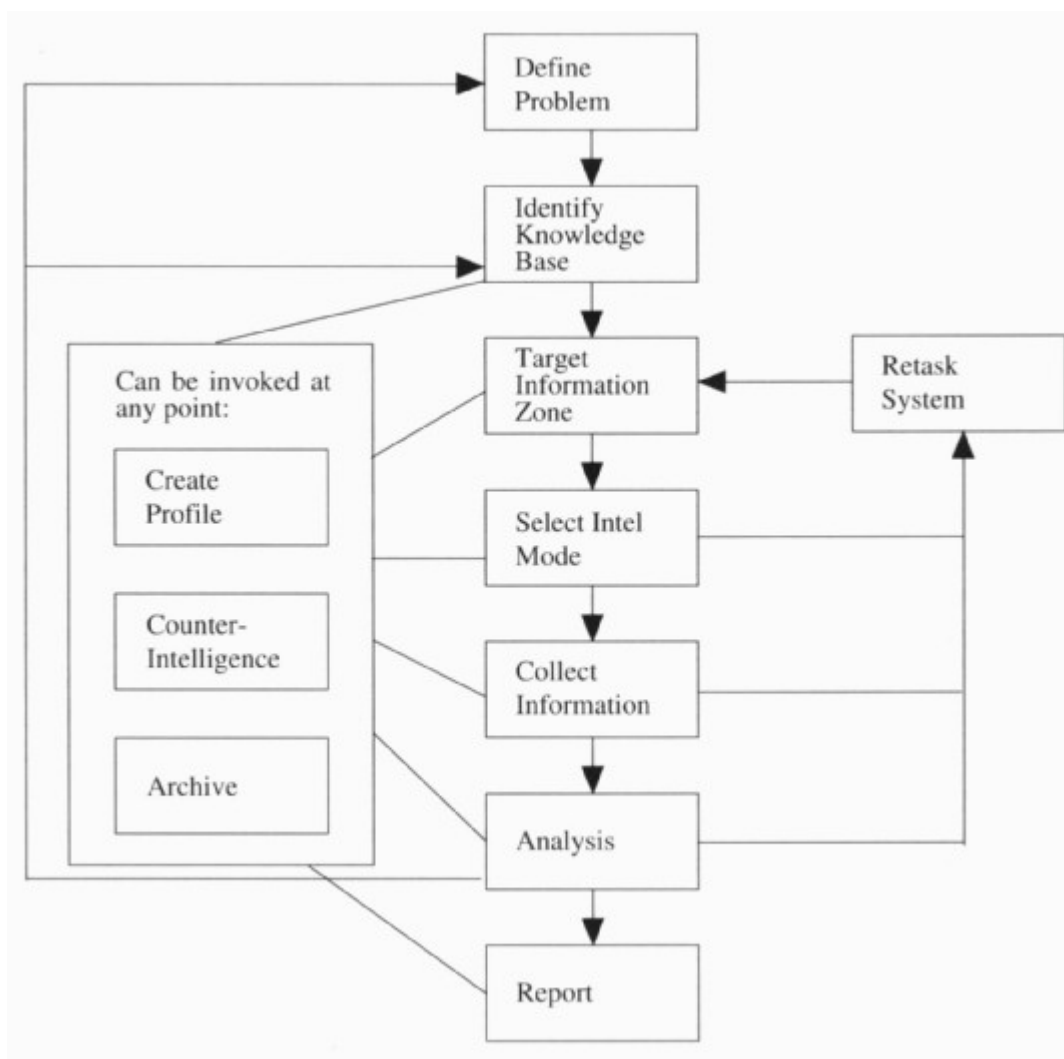
- Define problem
- Identify knowledge base
- Target location of information
- Select intelligence mode
- Collect information
- Analysis
- Report

The mini-phases are:

- Create profile
- Retask the system
- Counterintelligence
- Archive

The main phases are completed sequentially. The mini-phase "retask the system" is the method of performing a loop, starting back at "target location of information". Figure 1 is an overview of the process.

Figure 1
Overview of the Competitive Intelligence Process



For the purposes of this model, we identify two main actors: the intelligence provider and the intelligence user. The intelligence provider is the entity that directs and carries out the entire intelligence operation. Of course, in reality, the provider maybe an individual or an organization. In this model, the provider is referred to as an individual. The intelligence user is the entity that will receive the final intelligence product. It is the information needs and desires of this person/organization that determine the direction of the operation.

Define Problem

The intelligence provider receives a problem statement from the intelligence user. In general, this model distinguishes between information that the provider collects (raw information) and information that the provider has processed in some way. The problem statement is raw information. It may be in any form, from a conversation to a formal document.

The intelligence provider uses the problem statement to create a mission statement. The mission statement is the document that will guide the rest of the operation. It is made up of three parts: mission requirements, mission constraints, and user intentions. The mission requirements are the “what” of the intelligence operation. They are declarative sentences that define exactly what the intelligence provider is expected to do in order to achieve a successful intelligence operation.

The mission constraints are constraints on how the intelligence operation will be carried out and what form the final intelligence product will have. The mission constraints consist of time constraints and form constraints but may also include other constraints which do not fit well into these subtypes. The time constraints can be either ongoing or definite. Definite time constraints can be either relative or absolute. Ongoing time constraints must be defined in terms of time between deliveries of the intelligence product to the intelligence user. These times may also be relative or absolute.

Form constraints define what form the intelligence operation and the final intelligence product must have in order to ensure that the user accepts them. Form constraints include the set of ethics that are appropriate to the user, the intelligence system the user trusts, whether the product should be qualitative or quantitative, the format in which the final product should be presented, etc. (Barndt, 1994).

User intentions are the “Why?” of the intelligence operation. User intentions describe: why the intelligence user wants a certain piece of intelligence; why it is important to him/her/the organization. This information provides the intelligence provider with the context to recognize information that is significant to the intelligence operation.

There are many ways to determine the user intentions and form constraints. Performing a profile is one method that is discussed later.

The intelligence provider presents the intelligence user with the mission statement for his/her approval. Ideally, both the user and provider now understand one another so the intelligence user approves the mission statement. Further changes to the mission statement require interaction with the intelligence user. These changes should be rare and should be seen as profoundly affecting the course of the intelligence operation.

The purpose of defining the problem is to ensure that the intelligence provider understands the needs and desires of the intelligence user. The mission statement describes all of the criteria for a successful operation. Without a clear mission statement, the intelligence operation cannot succeed, by definition. The problem has been successfully defined when all of the parts of the mission statement have been completed and both the intelligence user and intelligence provider agree that the mission statement represents the shared understanding.

Identify Knowledge Base

After the mission statement has been successfully completed and before the intelligence provider begins collecting information, he should explicitly identify what is known and what needs to be known. The knowledge base is the structure that contains all of this information. The knowledge base has three parts: the information inventory, the assumption inventory, and the information requirements.

The information inventory is a list of succinct declarations that represent information relevant to the mission requirements. This information is regarded to be factual, pending the validity of its associated source. Each declaration is associated with one or more sources. We use a stipulative definition of information for this model. Something qualifies as information only if it is associated with a source and we intend to evaluate the worth of that source. An assumption is information the source of which we do not intend to evaluate.

Accordingly, the knowledge base contains an assumption inventory. Assumptions are relevant to an operation if they affect the intelligence provider's judgment. They should be explicitly identified. If they are to be used later in the analysis phase, the intelligence provider must decide whether or not the source will be evaluated.

The intelligence provider uses the mission requirements, the information inventory, and the assumption inventory to create the information requirements. The information requirements is a list of core questions and their associated sub-questions that target the information that the information inventory does not (or does not adequately) cover.

The goal of identifying the knowledge base is three-fold: to keep track of what is known and what needs to be known; to explicitly identify the information that affects the intelligence provider's judgment; to allow the intelligence provider to focus on answering a series of discrete, manageable questions.

Identifying the knowledge base is just the initial phase of setting up this knowledge base. The knowledge base will be iteratively refined and updated throughout the operation. Therefore, there are no real rules for judging the completion of this phase. The knowledge base is updated during the mini-phase "retask the system."

Information that is added to the knowledge base is no longer raw information. It has been considered and, perhaps, summarized by the intelligence provider. It is now part of the system.

Target Location of Information

Once the information requirements have been defined, the intelligence provider can begin to target the probable "location" of the information that will satisfy the information requirements. The term "location" needs some further explanation. Friedman et al. (1997) conceptualize information as residing in one or more "information zones". According to them, there are five information zones: electronically formatted (zone 1), paper formatted (zone 2), gossip (zone 3), gray zone (zone 4), and proprietary/secret (zone 5). The ease with which a piece of information may be retrieved, the time it will require to retrieve it, the resources required, the amount of "information emission" that retrieving it will produce, and the risk involved in retrieving it is determined by which particular zone a piece of information is located.

The goal of targeting the probable location of the information that will satisfy the information requirements is so that, in the next phase, the intelligence provider will be able to make a decision about which intelligence mode (explained later) to enter. The output of the "identify location of information" phase is an information requirements evaluation. The information requirements evaluation is the list of information requirements with each requirement labeled with the location in which it is likely to be found.

The intelligence provider uses knowledge of the nature of each of the information zones as well as past experiences to make the estimation. The value of creating the information requirements evaluation is that it allows the intelligence provider to more easily compare the information requirement from the current operation with the information requirements from previous operation in order to make a better estimation about the likely location of the required information.

This phase has been successfully completed when all of the information requirements have been mapped to one or

more information zones.

Select Intelligence-Gathering Mode

At this point, the intelligence provider has an idea of what information is needed and what kind of resources will be consumed in order to obtain this information. Now s/he must decide which information will be pursued. This decision is determined by the relationship between the available resources (time, money, equipment, and personnel), the estimated costs (estimated from information zone characterizations), and the value of the required information to the intelligence user.

Friedman et al. (1997) distinguish three modes of intelligence gathering: passive, semi-active, and active (explained further in the next section). In general, the intelligence provider should begin collection in the passive mode and only move to the semi-active and active modes if it is absolutely necessary. The information that is chosen to be pursued determines the intelligence-gathering mode. This decision should be made explicitly because each mode of intelligence gathering requires its own preparation.

The modes are differentiated by many factors. Many of these factors are inherent in the information zone with which they are associated. However, one of the most important factors that differentiates the modes is amount of information that is emitted by the intelligence provider in the course of collecting information.

The goal of this phase is to encourage the intelligence provider to methodically weigh the consequences, costs, and benefits of entering a particular mode of intelligence gathering. The issue of emitting information that can be detected by others is important because such an emission may destroy the competitive advantage that the intelligence product would have to the user and hence cause the operation to fail.

Collect Information

After weighing the alternatives, the intelligence provider begins to pursue the information that was targeted in the previous phase. In general, the intelligence provider should begin by mining information that is internal to the organization and then move to external sources (Friedman et al., 1997). The reasons for doing this are as follows. An organization already has a mechanism in place for gathering tons of information every day. Much of that information is gathered by people who already share the context of the organization. They are likely to be in a position to recognize significant information when they see it. Outside sources are expensive and do not share the context of the organization and, therefore, may be less productive.

The passive mode is the mode that should always be entered first. In this mode, the intelligence provider gathers information from zones 1 and 2. All of this information exists in the public domain and is free to anyone. Collecting this information involves very little time, skill, expense, or social interaction. Since it is freely accessible, collecting this information does not leave traces that could be detected by a competitor who is attempting to discover one's intentions.

A rule of thumb for determining when zones 1 and 2 have been exhausted is when the bibliographies of newly collected pieces of information contain references to the things that you have already gathered (Friedman et al., 1997). When it appears that the first two information zones have been exhausted, the system should be retasked (discussed later).

The semi-active mode should be entered only after zones 1 and 2 have been exhausted, there are still unsatisfied information requirements, and the value of the required information to the user exceeds the costs and risks associated with retrieving it. This mode involves collecting information from zones 3 and 4. It involves social interaction so appropriate measures such as performing a personality profile (discussed later), choosing a persona, identifying the appropriate contact method, and tracing social networks are required. For these reasons, the semi-active mode is slower, more expensive, and more difficult than the passive mode. The extra measures required by this mode also increase the amount of information that is emitted by the intelligence provider's actions.

For each instance of social interactions, a contact sheet should be created. The contact sheet consists of the name of the source, the task associated with this instance of interaction, the method of contact, the persona used, the results of the interaction (including the information retrieved), and further questions that resulted from this interaction. A rule of thumb for determining when zones 3 and 4 have been exhausted is when sources point to other sources that have already been examined.

The active mode involves gathering information from zone 5, the proprietary/secret zone. It requires the presence of the intelligence provider (or some agent, either mechanical or human) to detect an "information emission" (Friedman et al., 1997) from a target. This may require the organization of covert operations that includes agents and sub-agents to penetrate a target organization.

Retasking the System

Retasking the system is the mechanism for performing a loop in the larger "intelligence analysis" process. It can be

invoked at any point in the process. but it will most often occur after the completion of an iteration of the collection phase (in one mode or another). In this phase, the intelligence provider deliberately moves information, which has been collected and intermediately stored, into the knowledge base. This means that each piece of collected information has been processed or summarized in some way by the intelligence provider and it will now officially become part of the system. Information should be moved into the knowledge base if it affects the intelligence provider's judgment.

Once the information inventory and the assumption inventory have been updated, they should be mapped against the information requirements. The information requirements that are not covered by the two inventories are the gaps in the knowledge base. New information requirements should be added to these to comprise the updated information requirements list.

Analysis

In this phase, the intelligence provider generates possible solutions (hypotheses) to each of the mission requirements that are contained in the mission statement. Multiple hypotheses should be generated for each mission requirement. These hypotheses are then evaluated with regard to the evidence that is associated with them. Evidence is a role that is played by a piece of information from the knowledge base when it either supports or undermines a hypothesis.

According to Heuer (1999), the steps for evaluating multiple hypotheses are:

- Create a matrix of solutions and evidence.
- Create for/against lists for matching the information with hypotheses.
- Remove irrelevant information
- Evaluate the diagnosticity of evidence: for each piece of evidence, the intelligence provider counts the number of hypotheses it supports. The fewer hypotheses a piece of evidence supports, the more diagnostic it is.
- Remove evidence with no diagnostic value.
- Assess the likelihood of each hypothesis.
- Determine sensitivity (the sensitivity reference points to the pieces of evidence upon which a hypothesis depends).
- Identify key events.

The goal of this process is to leverage all of the information that the intelligence provider has collected and reduce the cognitive burden inherent in evaluating multiple hypotheses. After all of the hypotheses have been evaluated, the intelligence provider must decide on a recommendation. A recommendation is a concise declaration of what the intelligence provider believes is the solution to the mission requirements. Its purpose is to support the intelligence user's ability to make a decision rather than just supplying him/her with more information.

Report

Once all of the collection and analysis are complete, the intelligence provider's findings must be put in a form that provides value to the intelligence user. The intelligence product consists of the evidence, hypothesis evaluations, and recommendations. The report can also include the sub-phases: Create Profile (mini-analysis phase that is used to acquire certain kinds of information); Counterintelligence (intelligence provider identifies what are pieces of information that represent a competitive advantage to the intelligence user); Archive (consists of storing all of the information that makes up a certain intelligence operation according to the terms or features that characterize the operation).

Conclusions and Future Directions

A commitment to some model of an intelligence analysis process is necessary if we are to provide a computerized agent with knowledge about what steps make up the "correct" intelligence analysis process. Ideally, the intelligence provider would be able to customize this model, to a certain extent, in order to meet his/her needs. A tool called POINT (Problem Organization INtelligence Tool) has already been developed by the authors, using Visual Basic and MS-Access, to encode this model presented. This tool, and the encoded model, is now being tested and evaluated by analysts in the Federal Bureau of Investigation (FBI).

The most challenging issue concerns defining where within the intelligence analysis process the agent can provide valuable assistance to the intelligence provider. The terms and artifacts introduced by this model provide a means to characterize instances of intelligence operations. For example, as the intelligence provider creates the information requirements, an agent could compare terms contained in this artifact with terms contained in the information requirements evaluations of previous operations. By then identifying the information zones that were necessary to access and identifying the sources that yielded relevant information, the agent could suggest useful

advice about which intelligence mode to enter or what the cost of acquiring certain information might be. It could also take the initiative to retrieve information from sources that have been useful in the past, and hence, save the intelligence provider's time.

Another possibility is that, during the analysis phase, an agent could notify the user when a certain "likely" hypothesis relies on an "important" piece of evidence that is either an assumption or that comes from a source that has been unreliable in the past. We are exploring text mining and case-based reasoning further to possibly help in these areas.

Hopefully, more opportunities like those mentioned above will become apparent as we proceed to consider the intelligence analysis process as a collaboration between an intelligence provider, an agent, and the system (the model).

REFERENCES

- Aha, D. & Chang, L. W. (1996). "Cooperative Bayesian and Case-Based Reasoning for Solving Multi-Agent Planning Tasks", Technical Report, Navy Center for Applied Research in Artificial Intelligence, Naval Research Laboratory, Washington, DC.
- Atkins, K. (2000). "I'm off to see the wizard", ARIS Corporation, HYPERLINK <http://www.arrowsent.com/oratip/genwiz.htm> www.arrowsent.com/oratip/genwiz.htm.
- Barndt, W. (1994). *User Directed Competitive Intelligence: Closing the Gap Between Supply and Demand*, Quorum Books, Westport, Connecticut.
- Friedman, G. et al. (1997). *The Intelligence Edge: How to Profit in the Information Age*, Crown Publishers, New York.
- Gams, M. & Hribovsek, B. (1996). "Intelligent Personal Agent Interface for Operating Systems," *Applied Artificial Intelligence Journal*, 10(4), Taylor & Francis.
- Heuer, R. (1999). Psychology of Intelligence Analysis, Center for the Study of Intelligence, Central Intelligence Agency, [//www.odci.gov/csilbooks/19104/index.html](http://www.odci.gov/csilbooks/19104/index.html).
- Horvitz, F., Breese, J., Heckerman, D., Hovel, D., & Rommelse, K. (1997). "The Lumiere Project: Bayesian User Modeling for Inferring the Goals and Needs of Software Users", Microsoft Research, Redmond, WA.
- Intelligent User Interfaces Conference Proceedings* (1999). Association for Computing Machinery, New Orleans.
- Kahaner, L. (1996). *Competitive Intelligence*, Simon & Schuster, New York.
- Liehowitz, J. (Ed.) (1999). *The Knowledge Management Handbook*, CRC Press, Boca Raton, FL.
- Liebowitz, J. (2000). *Building Organizational Intelligence: A Knowledge Management Primer*, CRC Press, Boca Raton, FL.
- Liebowitz, J. & Beckman, T. (1998). *Knowledge Organizations: What Every Manager Should Know*, CRC Press, Boca Raton, FL.
- Meyer, H. (1987). *Real World Intelligence: Organized Information for Executives*, Weidenfeld & Nicolson, New York.
- Munoz-Avila, H., Hendler, J., & Aha, D. (1999). "Conversational Case-Based Planning," *Review of Applied Expert Systems*, Vol.5.
- O'Leary, D. & Selfridge, P. (1999). "Knowledge Management for Best Practices", *SIGART Intelligence*, ACM, Winter.
- Schreiber, C., Akkermans, H., Anjewierden, A., de Hoog, R., Shadbolt, N., van de Velde, W. & Wielinga, B. (2000). *Knowledge Engineering and Management: The CommonKADS Methodology*. MIT Press, Cambridge, MA.
- Soltysiak, S. & Crabtree, I. (1998). "Automatic Learning of User Profiles—Towards the Personalization of Agent Services," *British Telecom Technology Journal*, 16(3), July.