■ SHAWNEE VICKERY, Feature Editor, Eli Broad Graduate School of Management, Michigan State University

Intelligent Agents—Belief, Desire, and Intentions Framework Implementation Using LORA

by Fred Mills, Bowie State University; and Jagannathan Iyengar, North Carolina Central University

The development of new intelligent agents requires an interdisciplinary approach to programming. The initial challenge is to describe the desired agent behaviors and abilities without necessarily committing the agent development project to one particular programming language. What are the appropriate linguistic and logical tools for creating a top level, unambiguous, program independent and consistent description of the functions and behaviors of the agent? And how can that description then be translated easily into one of a number of program languages? This article provides a case study of the application of a simple Belief, Desire, and Intention (BDI) first order logic to a complex set of agent functions during the basic research stage of a community of intelligent nano-spacecraft. The research was conducted at NASA-GSFC (Greenbelt), Advanced Architecture Branch, during the summer of 2001. The simple examples of applied BDI logic presented here suggest broad application in agent software development.

Introduction

The demand for intelligent agent software is likely to grow as both public and private sector innovators seek to deploy adaptive, autonomous, information technologies to production, scheduling, resource management, office assistance, information collection, remote sensing, and other complex functions. Intelligent agents, also known as autonomous agents, are distinguished from other software programs by their ability to respond to a changing environment in pursuit of goals. One of the most critical stages of such intelligent agent development is the basic research that goes into determining how to translate client needs into agent software. Since programmers and operations managers have different expertise, there is often a linguistic gap between the functional language of the customer and the technical programming language of the agent developer. If the customer and agent developer do not speak a common language, both time and money may be lost in needless errors and misunderstandings. There is a need for a program-independent, unambiguous, logically consistent language for describing the desired computer agent's practical reasoning abilities and behaviors. Program independence gives the customer maximum flexibility in choosing a programming language to implement the desired functions. The avoidance of ambiguity allows the customer to say exactly what she means. And logical consistency ensures that only valid arguments will be generated by the agent's knowledge base, avoiding false beliefs about the world from being deduced from true beliefs.

A program's independent language for describing agent functions, practical reasoning, and behaviors solves what may be called the top-level description problem. The top level of description should be close enough to ordinary English to be interdisciplinary, yet rigorous enough to avoid ambiguity and inconsistency. The strategy of solving this problem borrows from first order logic, cognitive science, and recent work in the field of artificial intelligence, specifically autonomous agent theory (e.g., see the collection of articles in Muller, Wooldridge, and Jennings, 1997). The solution to the top level description problem suggested by some of the recent autonomous agent research is to combine the practical reasoning tools developed by psychologist Michael Bratman's belief, desire, and intention (BDI) framework (1987),

Fred Mills
is currently an assistant professor of philosophy at Bowie State University in Maryland. He has published articles in Journal of Applied Psychology *and* Journal of Mind and Behavior. *In 2001 he worked as a summer research fellow at Goddard Space Flight Center. He is an active member of American Philosophical Society and a member of the“Mind Project” under the auspices of University of Illinois-Chicago.*
fmills@bowiestate.edu

Jagannathan Iyengar
is currently a professor of management information sytems in the School of Business at North Carolina Central University. He publishes his research regularly in the areas of DSS, ES, and AI applications, with articles in Journal of Computer Information Systems, Journal of International Information Management, Journal of Information and Technology Management, Journal of Global Competitiveness, Competitiveness Review, Business Research Year Book, *and various proceedings. Dr. Iyengar has been awarded the Langley Research Fellow (2000), Goddard Fellow (2001, 2002), and Lawrence Livermore Fellow (1999).*
jiyengar@juno.com

with some of the tools of first order logic. This hybrid language is usually referred to as BDI logic. There are a number of BDI logics under development; here we employ the Logic of Rational Agents (LORA) developed by Michael Wooldridge (2000) to illustrate how BDI logic works in practice.

The authors faced the top-level description problem at NASA during the basic research stage of specifying the behaviors of a community of agents, in this case, the Autonomous Nano Technology Swarm (ANTS). ANTS will be designed to engage in practical reasoning and behaviors that implement the variety of functions necessary to explore the asteroid belt and communicate observations to earth.

Methodology

The first methodological consideration is to determine the level of detail that is appropriate in applying LORA as an agent development tool. The purpose of using BDI logic here is not to prove theorems; a sufficient body of deductive proofs has already been developed in standard first order logic texts. Nor is it to impress the client with obtuse descriptions of functions. It is rather to use just enough symbolism to economize and clarify the practical reasoning of the agents. The second methodological consideration is to announce our position on exactly what we mean to imply when employing psychological terms to describe computer agents. In what John Searle calls strong AI, conscious experience is attributed to intelligent computer agents (1984). In this paper we remain agnostic with regard to the debate over whether computer agents have real as opposed to "as if" intentionality (see Mills, 1998a, 1998b, for a detailed discussion). We discuss only those features of the debate that elucidate the advantages of using BDI logic as opposed to engineering or design idioms for top-level descriptions of communities of agents.

Part one of this paper provides theoretical justification for the use of Bratman's BDI framework for understanding agent behaviors. Part two provides an explanation of how first order logic can be combined with BDI and a dynamic component to account for agent decisions in time. Part three presents a practical problem of describing the behavior of a community of agents in a very complex scientific endeavor, in this case, the ANTS exploration of the asteroid belt. We believe the insights presented here have practical applications for the large variety of agents that will be developed in the near future.

Bratman's BDI Framework: The Intentional Stance

The intentional stance treats a sufficiently complex system as if it had mental states and engaged in practical reasoning. The term "intentionality," first used in empirical psychology by Franz Brentano (1973/1874), means directness towards an object. As John Searle points out (1984), one of the unique features of mental states is that they are always directed towards or about some object. If I perceive, I perceive something. If I believe, I believe that something is the case. If I intend something, I have some purpose in mind. We often use the intentional stance casually, as when we say the car does not want to start. But we do not really believe that a car has desires. Regardless of where one stands in the philosophical debate regarding whether computer agents have real or just "as if" intentionality, both sides generally agree that the intentional stance is an economic way to describe a complex system that engages in practical reasoning. Daniel Dennett (1978) points out that this economy of expression becomes clear when we contrast the intentional stance with the engineering stance. ANTS provides a good example of the economy of expression that results from opting for the intentional stance for top-level description of agent behaviors.

Imagine a community of nano-spacecraft setting out to explore the asteroid belt. At some point in the deployment of this community ANT #24 determines some activity is a priority and desires to execute that activity. At the programming level of description one would need to know the semantics and syntax of the programming language and read the code. At the even more detailed engineering level, one would then translate this higher level code into the impossibly complex machine code. By contrast, if one employs the intentional stance, the practical reasoning of Agent#24 becomes simple and concise: "Agent #24 desires to observe Asteroid Psyche, believes that the conditions are now optimal for such observation, and thus intends to make such observations." The next step is to imbed this intentional language in a first order logic to ensure the consistency of all descriptions of the ANTS behaviors.

Combining BDI with First Order Logic and Dynamics

First Order Logic

First order logic contains simple symbolic logic rules for combining propositions. For example, "Anna is a student and Anna is a biology major" can be represented by conjunction: $p \wedge q$. First order logic provides a generally accepted consistent logic. It is ideal for stating combinations and relations between beliefs, desires, and intentions about the world. It allows us to state just what the system knows without having to account for what it does not yet know (for a discussion, see Levesque and Lakemeyer, 2000). In order to represent the knowledge of the agent, some initial set of beliefs about the world are formulated. The inference rules of first order logic are used to ensure that the beliefs are consistent and that arguments are always valid. By following the basic rules of first order logic, new beliefs can be deduced from the current set of beliefs and acquired beliefs. (These acquired beliefs are based on new inputs from the environment.) In the case of ANTS, the initial beliefs are about known asteroid types, relative locations, shapes, rotations, mass, distribution, gravity, albedo and in some cases provisional classifications of asteroids [P.E. Clark, ANTS 04/05/01]. The inputs will originate in sensor instrument data and communications data.

An example of a representation in the ANTS knowledge base using some simplified LORA is:

- (differentiated)Psyche $\vee$ (undiff) Psyche [either, or]
- (differentiated)Psyche $\Rightarrow$ (PriorityObject) Psyche [if, then]
- (differentiated)Psyche {new knowledge!}
- (5000MiFrAnt#24)Psyche

Notice that first order logic is used here, without BDI, to represent the truism that asteroid Psyche is either differentiated or undifferentiated. If it is differentiated, it becomes a priority object of investigation that could trigger other actions. Assume that sensor information confirms that Psyche is differentiated. It then follows necessarily that Psyche becomes a priority object. The last formula states that ANT#24

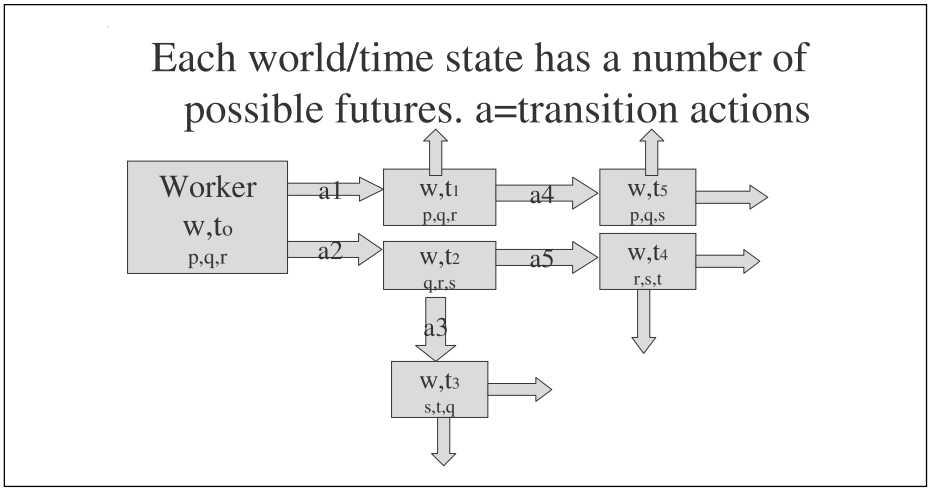

Figure 1: The world of an ANT (adapted from Wooldridge, 2000).

is in the proximity of Psyche. We can expect that, given the right combination of BDI, #24 will set as a goal the closer observation of Psyche and finally actuate behaviors towards the attainment of the goal. In order to begin the construction of a language that includes practical reasoning, Wooldridge adds a Belief component to represent epistemic states (2000). Thus we refine our description of ANT#24's knowledge by adding:

(Bel q PriorityObject(Psyche))

Here, agent #24 (represented by the symbol "q") believes that pysche is a priority object. This belief can be updated or modified given additional information from the sensors or new communications. In order to complete Bratman's framework for practical reasoning, we now add the intentional and emotive components: intentions and desires. Consider a desire as an enduring goal. An agent may have a number of desires, only some of which are realizable at any given time. An intention is a pro-attitude, that is, it is a movement toward an immediate objective until that objective is fulfilled or some new event changes the current intention (Wooldridge, 2000). This distinction between intentions and desires helps us to explain how an agent may change course in response to environmental variables that change through time. Here is an example of epistemic and emotive states combined, using LORA:

- (Des ruler (Bel worker#24 PriorityObject (Psyche)))

Here the ruler agent desires that ANT#24 believes that psyche is a priority object. This mental state is important because it may trigger an action by ANT#24, which has among its desires a desire to pursue priorities set by the ruler ANT.

Dynamic (Temporal) Component

Now that first order logic and BDI have been combined there is one more critical step required to complete a top-level description of autonomous agency: dynamics. In order to represent practical reasoning and resulting behaviors through time, a schema of relevant features of the present state of affairs is constructed. Since not every feature of the universe can be represented, the state of affairs, that is, the world of the agent, contains beliefs about only those features of the universe relevant to the functions of the agent. Since some agent technologies will be able to learn from their experience, in some architecture the relevant states of affairs can expand and change. For the present purposes, humans will define the relevant state of affairs. Each node in the schema will branch out into possible worlds. These possible worlds are alternative paths that the agent may select, depending on its current beliefs, desires, and intentions. The map in Figure 1 illustrates a world time schema, where w,t=world,time; p,q,r,s,t are beliefs; a=transition action; branches represent possible worlds (see Wooldridge, 2000, for a more detailed description of agent behaviors through time; Figure 1 is adapted from pp. 56, 63, 94)

Specify functions

We are now ready to specify the functions of various agents in the community of agents and then apply LORA to describe their behaviors.

General ANT community functions

- mapping and close up imaging of Asteroids using multi-spectral band coverage [Clark, Memo 04/05/01]

Worker ANT functions

- Communications, resource management, navigation, local status (housekeeping), local conflict resolutions Science data acquisition processing [Curtis et al. 2000, 3]

Ruler ANT functions

- Plan assignments for worker ANTS
- Maintain shared SWARM statistics
- Resource management [see S. Curtis, et al., 2000, 3]

The ANTS is conceptualized here as a community of autonomous rational agents whose behaviors are generated by a knowledge base (KB), a set of goals, inference procedures, and percepts. Although each worker ANT is autonomous in terms of its own function, it is subordinate to the function of the ruler. The intentions of the ruler are passed on to the workers through the messengers and the workers actuate the plans that achieve the goal. Each worker has as its permanent goal the appropriate and timely collection or discovery of data from the target type of objects, the maintenance of health and safety, and the timely communication of data to the messengers, who in turn report to the ruler and to earth. The workers' goals are sub-goals of the ruler's goal and the worker's plan at any given state of affairs is a sub-plan of the ruler's plan.

Example of a LORA Specification

In the following scenario, the Ruler has received information about an opportunity to view Psyche under ideal conditions and there is a group of worker ants in the neighborhood of Asteroid Psyche (A). The Ruler forms an **intention** to attain a goal as a result of a deduction that employs beliefs in its KB, acquired beliefs derived from current percepts, and its belief derived from communications with all ANTS. There is a potential for a group of workers and allied messengers to achieve the goal. This is exactly the sort of scenario supported by BDI type logic! For simplicity we leave out temporal considerations, quantification, and proofs, and goal/sub-goal relations. We seek to illustrate here the usefulness of BDI logic to model the practical reasoning or "mental states" of cooperating agents.

A = constant for Psyche

PfC = perception of the potential for cooperation

i = ruler agent

j = messenger agent

g = group of worker agents, each with specialized observation instruments

α = action

φ = goal to study Asteroid Psyche

achvs = attain through an action

Bel = Belief

Int = Intention

Des = Desire

π = Set of plans to be actuated (executed) by each worker to attain φ

ψ = Preconditions for φ to become the next goal (Science agenda from humans, combination of percepts in relation to KB).

J-attempt = joint attempt.

Assumption

The Ruler has formed an intention to achieve the goal of collecting data about Asteroid Psyche. We do not here represent the deliberations that lead to this intention. We begin with the process of mobilizing the workers to achieve the goal and a description of the mental state of the community of agents. The Ruler forms an intention to attain a goal (to study Psyche) as a result of ψ having been met. ψ, in this case, is a combination of the ruler's KB, mission priorities, current percepts, its belief that the goal has not yet been attained, its belief that the ruler itself cannot or does not desire to achieve the goal by itself, and its belief that there is a potential for a group of workers and allied messengers to achieve the goal. The ruler will therefore intend not only the goal, but that the messenger intends the goal and that the group of workers intends the goal and finally achieves the goal. This is exactly the sort of scenario supported by LORA and other BDI type logic!

Commentary

$(\text{Int}\ i\ \varphi) \wedge (\text{PfCi}\ \varphi) \wedge (\text{Int}\ i\ (\text{Int}\ j\ \varphi\ (\text{Int}\ g\ \varphi)))$

The ruler intends the goal of studying the Asteroid Psyche and acknowledges the potential for cooperation among the workers to being about the goal and intends that the messenger intends that the group of workers intends the goal.

$\text{Int}\ i\ (\text{Achvs}\ g\ \alpha\ \pi\ \varphi)$

The ruler intends that each member of the group following its share of the plan can achieve the goal through the group's actions.

$\{\text{Inform}\ i\ j\ \alpha\ \pi\ \varphi\} \wedge \{\text{Request Th}\ i\ \alpha$
$\{\text{Inform}\ j\ g\ \alpha\ \pi\ \varphi\}\}$

$\{\text{Agree}\ j\ i\ \alpha\ \alpha'\};\ \alpha'$

Here the α is the informing action and α' is the action that j agrees to perform.

$\{\text{Inform}\ j\ g\ \alpha\ \pi\varphi\} \wedge \{\text{Requests Th}\ j\ g\ \alpha\ (\text{Happens}\ \alpha')\}$

It is assumed that, unless the workers receive percepts, in relation to their KB, that requires that the goal be re-evaluated, the group will believe that it can execute the plan to achieve the goal.

$(\text{M-Bel}\ g\ \neg\ \varphi) \wedge (\text{M-Des}(\text{Achvs}\ \alpha\ \pi\ \varphi)) \wedge$
$(\text{M-Int}\ g\ (\text{Achvs}\ \alpha;\ \varphi?)) = \{\text{J-Attempt}\ g\ \alpha\ \varphi)$

Conclusion

The advantage of using a language like LORA is that we can always add new functions to an agent and make agent behaviors depend on previous knowledge (the set of beliefs in KB and new percepts). We can also add to list of possible behaviors by adding new beliefs. Most important, we can very economically describe the behavior of the system at a "high" level without employing program language.

The ANTS concept is indebted to the work of Mark Campbell, et al. work in "Intelligent Satellite Teams for Space Systems." As an example of an Intelligent Satellite Team (IST), the report considered "a large number of satellites orbiting the Sun using optical and infrared sensors to look for asteroids." These satellites would share data, track, and collect information about interesting asteroids. The report mentions the important role of practical reasoning in IST: "Reasoning in an agent allows the agent to reach an intelligent conclusion or decision about a given set of data." The report notes that this practical reasoning would give the agent more adaptability and autonomy. But there is no attempt to develop a top level language for describing this type of reasoning.

In another foundational paper, Steven Curtis, et al. (2000) provide a description of the mission, system architectures, and ANT functions, including the type of instruments ANTS will employ to collect data about asteroids. There are also descriptions of how ANTS might coordinate their functions. The operational scenarios, however, do not attempt to represent the state of knowledge of each ANT, and therefore do not reveal the anatomy of practical reasoning. Again, a language of ANT practical reasoning would be required to carry out such an objective. Truszkowski, Zoch, and Smith point out that a subsystem agent must be able to exercise some "reasoning process" in order to make use of external resources in new ways. The first stage in realizing spacecraft constellation autonomy is "agent development." We suggest that this stage requires an economic way to describe the practical reasoning of the agents.

References

Bratman, M. (1987). *Intentions, plans, and practical reason*. Cambridge: Harvard University Press.

Brentano, F. (1973). P*sychology from an empirical standpoint*. New York: Routledge. (Original work published 1874).

Campbell, M. (principle investigator). (1999). *Intelligent satellite teams for space systems*. Final report for the NASA Institute for Advanced Concepts Program. Seattle, Washington.

See RESEARCH, page 25

Metters is the shortest of the three and covers the fewest topics, but some of the coverage is better than the other texts. Topics covered in more depth than the others are entire chapters on: Managing "Experiences," Data Envelopment Analysis, Yield Management, and Scoring Systems (e.g., can be used for rating customers in a lending operation). Also, the discussions of yield management, data envelopment analysis and especially inventory management are better at showing the service context.

Traditional topics not covered include forecasting, scheduling, and quantitative methods for quality management (there is a qualitative discussion of quality management). This suggests that Metters is likely more appropriate for an advanced- or MBA-level service operations course. One other thing not present is a section of "solved problems" in the quantitative chapters, although this is partly mitigated by the end-of-chapter problems with good instructor solutions, some of which are complete with extensive explanations.

Another aspect that suggests this text can be appropriate for a more advanced class is that the end-of-chapter problems and cases at times go beyond the chapter discussion (more than is typical in many texts), "pushing" students and forcing them to think in different ways.

Concluding Thoughts

Although strengths and weaknesses have been outlined above, all three can be used for a service course at any level, partly due to the good instructor support materials. Personally, I would not attempt to cover all the material in any of the texts in a twelve- or thirteen-week semester, rather choosing to have richer classroom discussion on fewer topics. If the semester is fifteen weeks, it would likely be possible to cover Davis or Metters in a single semester by spending less time on the nonquantitative chapters, especially if no classes are allocated for non-text material (such as student project presentations).

One thing I would like to see in the future (I can dream, right?) is a Canadian edition of a service operations text, although I'm not sure if the market is large enough.

On another personal note, I have recently started teaching the undergraduate introductory operations course with a stronger spreadsheet emphasis. One thing that all three texts lack is good spreadsheet support (e.g., spreadsheet examples including formulas and model-building in the text and also problems and solutions that "gently" push the students to higher model-building levels). Davis is likely the best of the three for this, but even that text does not have spreadsheet coverage in all quantitative chapters. Metters includes quite a few spreadsheets in the instructor materials as answers to the end-of-chapter-problems and the cases, but does not discuss spreadsheets much in the text.

As already stated, it's great to have a number of service operations texts available and all the authors are to be commended for their unique approaches to this area. ■

Peter T. Ittig, Feature Editor
College of Management
University of Massachusetts
Boston, MA 02125-3393
http://www.faculty.umb.edu/peter_ittig/

RESEARCH, from page 19

Curtis, S., et al. (2000). *ANTS (Autonomous Nano Technology Swarm): An artificial intelligent approach to asteroid belt resource exploration*. American Institute of Aeronautics and Astronautics.

Dennett, D. (1978). *Brainstorms*. Cambridge, Massachusetts: The MIT Press/A Bradford Book.

Iyengar, J. (2001). *Intelligent Software Agents-Human Agent Interface Agent Interaction*. NASA-ASEE 2001 Research Paper.

Levesque, H., & Lakemeyer, G. (2000). *The logic of knowledge bases*. Cambridge, Mass.: MIT Press.

Mills, F. (1998a). The easy and hard problems of consciousness: A Cartesian perspective. *The Journal of Mind and Behavior, 2*, 119-140.

Mills, F. (1998b). A critique of the concept of computer moral agency. *Humanities and Technology Review, 17*, 26-39.

Muller, J., Wooldridge, M., & Jennings, N. (1997). *Intelligent agents III: Agent theories, architectures, and languages*. ECAI'96 Workshop (ATAL) Budapest, Hungary, August 12-13, 1996, Proceedings. Berlin: Springer-Verlag.

Truszkowski, W., Zoch, D., & Smith, D. (1999). *Autonomy for constellations*. Research paper. GSFC, NASA.

Wooldridge, M. (2000). *Reasoning about rational agents*. Cambridge, Massachusetts: MIT Press. ■

Shawnee Vickery, Feature Editor
Broad Graduate School of Management
Michigan State University
vickery@msu.edu

DSI Online Membership Directory Is Available to All Institute Members

To access the DSI online membership directory, point your browser to:

http://decisionsciences.org/directory/auth.asp

Log in using your email address and membership number as password. If you do not remember your membership number, you can have the system send your password to you from the authentication page. If you have not registered an email address, contact us at **dsi@gsu.edu** with your membership number and email address, and we will update your profile for you. Once you are logged in, you can also update your own profile. ■