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

Intelligent Software Agents for Future Space Explorers: A Conceptual Framework

by Jagannathan Iyengar, University of Wisconsin-Whitewater; Walt Truszkowski and Pam Clark, NASA Goddard Space Flight Center

Jagan Iyengar
is currently associate professor of management information systems at the University of Wisconsin-Whitewater, MCS Group and Department of Management. He publishes regularly in the areas of DSS, ES, and AI applications, and participates in Decision Sciences Institute annual meetings by presenting papers and discussing papers in DSS/ES/AI areas. Dr. Iyengar has published in Journal of Computer Information Systems, Journal of International Information Systems, Journal of Global Competitiveness, Competitiveness Review, Business Research Year Book, *and proceedings of the Decision Sciences Institute and its regions. Dr Iyengar received Senior Faculty research fellowships at the NASA-Goddard Space Flight Center during 2001 and 2002, NASA Langley Fellowship in 2000, and Lawrence Livermore Fellowship in 1999.*

iyengarj@uww.edu

Walt Truszkowski
is a Project Leader from NASA Goddard, code 588. Mr. Truszkowski is the coordinator for Agent Community Projects and weekly seminars at Goddard. He also organized Workshop on Rational Agent Community during 2000, 2001, and 2002 at Goddard Space Flight center. Mr. Truszkowski also handles Summer Student fellowships at NASA in Agent Based research. Mr Truszkowski is also a leader in Semantic WEB community at University of Maryland, College Park. Recently, Mr Truszkowski published four articles along with Dr. J.V. Iyengar in the areas of Rational Agents Behaviour and HUman Agent Interaction, Human Agent Interface, and Intelligent Software Agents as Space Explorers for ANTS Mission.

Pam Clark
is research scientist at NASA-Goddard, code 691. Dr. Clark participates in Agent Community research projects at NASA. She is also involved with scientific mission details in Asteroids related studies. Dr. Clark also teaches astrophysics at University system of Maryland periodically.

NASA has pioneered the use of Intelligent Software Agents (ISA) in sophisticated open-ended human endeavors, namely, missions of remote exploration of the Earth and Solar System. As onboard data storage capacity, number of operating spacecraft, and the complexity of mission technical and scientific goals have increased, greater cost constraints have led to fewer resources and the need to operate spacecraft in space with less direct human interaction. Thus, the role of ISAs in space inevitably continues to evolve. Autonomous operation of spacecraft systems has been a requirement and a reality for the past decade. Often, autonomous operation has successfully allowed spacecraft to 'safe' themselves when human error or environmental conditions compromise spacecraft safety. Currently, planning stage missions such as TRMM will require teams of spacecraft to cooperate at a tactical level to achieve predefined specific monitoring goals: selecting and scheduling measurements to be made by appropriate instruments to characterize rapidly unfolding real-time events on a routine basis. The next level of development, currently under discussion, is in the use of ISAs at a strategic level, to explore targets associated with a large class of objects, with only infrequent human interaction possible. Supervised clusters of spacecraft will operate simultaneously within a broadly defined framework of goals (e.g., determine the taxonomic relationships between mainbelt asteroids and implications for origin of asteroids), to select targets for study from among a large number of potential candidates, to collect only enough data to characterize each target's relationship to the entire set of objects. What particular traits of intelligence are most valuable for agents to possess in this environment (see Figure 1)? What type of ISA architecture(s) would be optimal for organizing such a community of agents? What combinations of modalities (see Figure 2) in decision level (operational, tactical, or strategic), reasoning (inductive or deductive), cognitive style (intuitive or sensing), nature of activity (data-based, model-based) are best suited for the level of complexity (well-defined task, open-ended problem) and organizational stratum (worker, supervisor, leader) in this environment? How well should input (goals, models, databases) be predefined (and occasionally redefined by human interaction) to allow learning capability to evolve within the agent community? In this paper we will conceptually model the successful remote exploration (the open-ended problem) environment in terms of combinations of modalities and ISA architecture which will allow optimal achievement of goals, and construct a scheme for testing our models.

ISA Relevant Characteristics, Environmental Modalities, and Architectural Models

Characteristics

We use the ASAP acronym of a way of encapsulating the most relevant subset of intelligence characteristics from among many that have been discussed by previous authors. In a complex problem-solving environment where communication between agents may be sporadic over great distances and contact with human agents rare, **autonomy**, the ability to perform with little external intervention, is crucial and separates successful ISAs from programs. On the other hand, when communication does occur, it must be beneficial in achieving the larger goals; thus, another key requirement is **sociability**, the need to be a team player in a big scheme, resolve conflicts, follow directions in the best interest of the group, to collaborate, cooperate, and coordinate as necessary. A high degree of **adaptability** will be required: the ability to figure out a way to obtain information more optimally and to achieve more relevant goals as the environment changes and/or new information becomes available, in other words, to learn and evolve over time. Finally, **persistence**, the will to survive, and, if necessary, find a way to come back online despite setbacks, in order to achieve goals, is as crucial for non-human agent explorers as it is for human explorers.

*A*utonomy:	ability to operate in stand-alone mode indefinitely
*S*ociability:	need to coordinate/cooperate/collaborate for the community goals
*A*daptability:	ability to evolve in operations/tactics/strategies with new input
*P*ersistence:	ability to survive and achieve goals despite obstacles, loss, and change

Figure 1: Most relevant characteristics of intelligence in remote, open-ended problem-solving agent communities: ASAP.

	Level 1	**Level 2**	**Level 3**
Decision:	Operations	Tactics	Strategy
Organization:	Worker	Supervisor	Leader
Intelligence:	Rule–oriented	Cooperation	Collaboration
Complexity:	Data-gathering	Data-products	Pattern Identification
Input:	Database	Model	Lofty Goal
Activity:	Data-based	Data+Model	Model-based
Reasoning:	Inductive	Inductive/Deductive	Deductive
Cognition:	Sensing	Sensing/Intuitive	Intuitive

Figure 2: Modality End—Members for an open-ended problem-solving environment.

Learning/Knowledge	**Architecture**	**Application**
Static/TopDown	Expert Systems	Instrument calibration
Limited/Coordinated	Constrained Genetic Algorithms	Scheduling observations
	Relational Databases	Data Mapping
	Artificial Neural Network	Derivation of data products
Evolving/Distributed	Fuzzy Logic	Select targets

Figure 3: Relevant ISA architectures in context of space explorers.

Environmental Modalities

Each modality defines one aspect of the agents' activity and operating environment and the spectra of behaviors, from least to most complex, the agents could have in response. A particular activity may call for operations (perform well-defined task, e.g., turn on instrument), tactics (coordinating operations, e.g., schedule data-taking activities of a number of agents), and strategy (review results and request moving on to another target) which requires degrees of intelligence ranging from simply following instructions (rule-oriented for well-defined tasks), to cooperating with others or coordinating others, to collaboration with many agents (distributed intelligence). The activity would be performed at the worker (operator), supervisor (tactician), and leader (strategist) level. If defined in terms of taking measurements with an instrument, activities would range, in terms of complexity, from data-gathering, to data-product generation, to pattern identification. The input at the corresponding levels would be from a database, model, or lofty goal. The equivalent nature of activity would range from data-based to model-based, types of reasoning from inductive to deductive, and nature of cognition from sensing to intuitive.

ISA Architectures

A variety of ISA architectures, as illustrated in Figure 3, have already been developed for real-world applications, and are often used in combination. An Expert System (ES) uses a top-down, rule-based, object-oriented, hierarchical approach, and is particularly well-suited to solving well-defined problems in a controlled environment. Control and intelligence are not evenly distributed, and do not evolve over time, but reside at the top of the hierarchical chain of command with the 'expert'. An example is the system developed (Hanson & Brekke, 1988) to collect data in order to maintain networks or workstations at NASA. The Java Application Development Environment (JADE) software development tool uses this approach. Relational Databases (RD) show relationships between input datasets on the basis of their characteristics within predefined categories (as identified by keywords). Such approaches are used in spatial mapping, where values for a range of input datasets are associated with given spatial bins, and can be mapped in 2- or 3-dimensional space. Keyword searching, or data mining, both used in e-Commerce applications, are also based on this approach. Constrained Genetic algorithms (CGA) typically solve problems in scheduling by treating input as time series strings, equivalent to gene sequences, and systematically dividing and recombining portions of these sequences until a solution appears according to predefined constraints or rules (e.g., downloading data from multiple spacecraft agent data buffers at varying bit rates according to availability of receiver agents of varying capabilities). Artificial Neural Networks (ANN) are typically used for pattern recognition. ANN software 'learns' to recognize systematic features with a carefully selected training subset of the entire dataset, is tested on a randomly selected testing subset, and then is deployed to identify systematic features for the entire dataset. Commercial applications include recognition of credit card fraud. A typical scientific application would be learning to recognize signal from background in collected spectra. Fuzzy Logic (FIDE, Fuzzy Input Decision Environment) is applicable where a problem is not well-defined, but can be determined in terms of probabilities that necessary fuzzy criteria are met for membership in a 'fuzzy set.' The logic in antilock braking systems is a commercial application of fuzzy logic. Such an approach is particularly applicable to open-ended problem solving, where a specific parameters are evaluated in terms of their best 'fit' to fuzzy sets defined by a pre-existing model. As Figure 3 indicates, the suitability of these architectures is dependent on the amount of learning which will be required in the application environment. Evolution of agent behavior can be seen from top to bottom in the figure as (1) following rules with knowledge and control flowing from the top down (expert systems), to (2) changing in typical behavioral responses as database becomes more inclusive with coordination flowing from the top and cooperation from the bottom (artificial neural net) to (3) collaboration of team members with various inputs leading to distributed intelligence among a community of agents or multi-agent systems (fuzzy logic).

System	Worker	Supervisor	Leader
Propulsion system	**	*	*
Guidance/ Navigation	**	*	*
Attitude Determination/Control	**	*	*
Communications	*	**	*
Command/Data Handling	*	**	*
Power	*	*	*
Structures/Mechanisms/Thermal	*	*	*
Scientific Instrument	*		
Artificial Intelligence	*	**	***

Figure 4: Primary spacecraft subsystems.

Instrument	Optimal Maneuver	Viewing Requirement
	Bulk Property –> Map	
Imaging Spectrometer	Flyby–>Hover–>Orbit	Full to partial illumination
Ranging Device	Flyby–>Hover–>Orbit	Nadir-point
Radio Science	Flyby–>Hover–>Orbit	Close
Near IR Spectrometer	Hover–>multi-Hover	Full illumination, nadir-point
X-ray Spectrometer	Hover–>multi-Hover	Close, Full illumination, nadir-point
Gamma-ray Spectrometer	Hover–>multi-Hover	Close, long integration times, boom
Magnetometer	Hover–>multi-Hover	Close, boom

Figure 5: Probable scientific instruments and operational requirements.

Present Space Applications

The nature of use in context of space exploration applications, specific missions to application spacecraft systems agents (SOHO) and the more recent Earth (TRMM) mission examples would be good for building the framework for Intelligent Software Agents (ISA) representations.

Future Space Explorers

In our conceptual framework (see Figure 4), space explorer agents will operate as teams of nanospacecraft (10 kg), each spacecraft bearing at least one scientific instrument on deep space missions. Agents will be required to do more than survive in a challenging environment and with minimal human intervention. Such swarms, communities of multi-agent systems, will be designed to explore a whole population of related objects. The specific overarching goal will be to return data products which are based on high quality measurements from targets selected in a way which accurately represents the population being studied. A priori knowledge will include crudely calculated trajectories of the larger objects in the population, derived from ground-based measurements, as well as current paradigms for classification based on the smaller portion of the objects for which ground-based spectral measurements are available. Potential candidates for such observations include the satellite systems (Uranus, Neptune…), asteroid or comet populations (mainbelt, Earth crossers…).

Our present design calls for a community of space exploration agents. The heart of the community is the three types of autonomous spacecraft 'specialists' organized in a minimal hierarchy. Spacecraft subsystems are listed in Figure 4, with asterisks indicating relative size and importance for each type of specialist. Artificial intelligence heuristic systems and system interfaces involve every spacecraft function. Workers, nanospacecraft each bearing one of the single instruments listed in Figure 5, operate in data (sensing) mode to perform data-gathering and instrument maintenance. Supervisor/ Messenger agents provide tactical assignment and analysis of worker operations and thus must be able to perform in data (sensing), model (intuitive), or combined mode. Leader/Chief Knowledge Officers provide strategic evaluation and interpretation of data, as well as long-term planning, such as evolving selection criteria for targets, and thus perform primarily in model (intuitive) mode. The worker to supervisor to leader ratio is approximately 100:10:1. The Leader may have intermittent contact with the human 'agent.' The targets of study can also be considered agents.

Activity	Agent	Style	ISA Approach
I. Acquiring Asteroid Agent	Chief	Strategy N	FIDE
Or Leaving Asteroid	Supervisor	Tactics N/S	
(target valid/non-valid)	Workers	Operations S	
II. Scheduling observations	Supervisor	Tactics N/S	GCA
III. Generating Raw Data	Workers generate	Operations S	ES
	Uncalibrated:		
	VI image		
	NIS,XRS,GRS spectra		
	RD,RS time series		
IV. Deriving Measurements	Supervisor with	Tactics N/S	ANN, RD
	Workers generate		Look for 'signal in
	Calibrated, normalized		discrete spectral
	VI image, light curve		regions defined
	Albedo, color		via prelaunch
	NIS band spectra		training by humans
	XRS,GRS line spectra		
	RD topography profiles		
V. Analyzing measurements	Leader w/ Supervisor	Strategy N	FIDE, ANN
	3D shape model		3D shape model
	Closest Asteroid Type		done early by
	Closest Meteorite Type		imager to allow
	Geochemical Analysis:		spatial mapping
	Mg, Al, Si, Ca, Fe, O, K, S		
	Mineralogical Analysis:		
	pyroxene, olivine		
	Bulk estimate and		
	Mapped abundances		
	Albedo/Color/morphology		
	texture unit 'maps'		
VI. Interpreting Results	Human w/ Leader	Strategy N	FIDE
	(Systematics for asteroid: Nature of variations in surface (modification thru geological processes) and interior (differentiation thru geochemical processes) leading to understanding of origin w/in asteroid belt		
	Systematics for asteroid belt: trends and relationships between composition, size, surface roughness, dynamics leading to understanding of origin w/in early solar system)		

Figure 6: Exploration parameter matrix for ANTS—A proposed asteriod mission.

Target Selection, Scheduling and Acquisition, and Validity

Targeting activities are essential for the exploration process and thus will be done with collaboration among all spacecraft agents, because such activities involves strategy, tactics, and operations. Summarization of the nature of this and other explorer activities is shown in sequence in Figure 6. At all times for this and all subsequent tasks, spacecraft agents must know their positions relative to the closest team members and supervisor, as well as the position of potential targets and the sun. Instrument bearers must also know where their instruments are pointing. Distributed a priori knowledge consists of a model of potential target trajectories and other relevant target data where available. Further sensor input on target positions gained during the mission are used to update the original model.

1. Recommendations for target selection will be generated at the strategic level, probably with some version of fuzzy logic, based on current state of knowledge of the target population. The spectral classification of the target, in some cases available a priori or in other cases measured early in the target encounter, will determine the level of study. For more 'primitive,' undifferentiated, relatively homogeneous targets, determination of the bulk property (e.g., average iron abundance) through single measurements would be adequate. Distribution of a property in the form of maps with some spatial resolution would be called for in the case of more evolved, differentiated, heterogeneous targets. The level of study, which has implications for optimal maneuvers during data-taking operations (Figure 5), is based primarily on two parameters: coverage (extent of maneuvering, e.g., number of hovers, at target) and data quality (length of time at target). Fuzzy logic is also used at the strategic level to determine when a target should be 'deselected' or abandoned. Optimal target parameters evolve as the state of knowledge evolves. The validity of a particular target as a candidate for data-taking activities is reevaluated at after each download of an instrument data buffer. When sufficient data quality (signal to noise ratio) and coverage (complexity of maneuvers) have been achieved, according to the current evolving strategy or operational limits, a target is no 'deselected' for a particular instrument.

2. Scheduling of observations, as well as the downloading of observed data buffers, is done at the tactical, or supervisory, level based on the availability and condition of targets and workers in communication range. Feedback from workers' measurements at this level which indicate a morphologically complex surface (indicating geological evolution) would generate interest in a higher level of study and result in the scheduling of more observations. The capability of worker propulsion systems would be another variable to consider. Scheduling involves use of target and spacecraft models as well as current data. AI architectures which are optimal for scheduling, such as Constrained Genetic Algorithms, could be used here, combined with fuzzy logic for identifying constraints.

3. Typical instrument operational requirements are indicated in Figure 5. Maneuvers used to acquire data from the target

will vary with the instrument as well as with the nature of the target. The more complex targets are the more complex the maneuvers which will be required. Implicit, but not stated, is the requirement for spectrometers that the respective fields of view be filled during data acquisition, and that provisions be made for periodically looking away from the target to acquire background data. Note that the optimal operational scenarios vary substantially from instrument to instrument. Instruments are listed in rough sequence, in terms of stringency of operational requirements, by distance (furthest to closest) and time (least to most) requirements in Figure 5. Most likely, the visible imaging spectrometer will be used to ascertain the exact location and create a rough model of the target prior to the arrival of other instruments. The required viewing conditions for the target vary considerably, and might well affect the choice of targets. Near Infrared and X-ray spectrometers, which require full illumination and nadir-pointing, would not be sent to close target if it were poorly illuminated. Target acquisition would be carried out by each worker at the operational level, on the basis of optimal viewing conditions. As long as these are well-defined, heuristic, object-oriented programming could be employed in this process.

Observation and Analysis

Analysis of observations requires many steps, summarized in matrix form by agent, data type, activity and AI style, in Figure 6.

1. Generating raw data: According to predefined specific constraints, instruments are calibrated, pointed at the target optimally for viewing constraints, and measurements are recorded, all as operations by single instrument spacecraft, or workers. Forms of raw (uncalibrated) data are listed by instrument in Figure 6. Heuristic, Object-oriented approaches would probably by appropriate for this inductive step.

2. Deriving measurements. Supervisor interacts with workers to derive fully calibrated, normalized (for source variation and viewing geometry), clean (noise or background removed) measurements, listed in Figure 6, in this step. Calibration (for instrument variations) will probably occur at the instrument/worker level. Noise removal and normalization steps will require a combination of inductive and deductive reasoning, more sophisticated modeling and handling of data from other from many sources, such as other instruments, and thus will probably be performed at supervisor level. Use of artificial neural net approaches could be used to recognize signal in discrete spectral areas defined and taught before launch (to separate real signal from sources of noise).

3. Analyzing Data: The leader interacts with supervisors to analyze derived measurements to create data 'products', a process which is both inductive and deductive. A large range of derived input data are combined to create models for parameters of composition, shape, surface complexity, while models for planetary object and surface modification are used to assess the meaning of input data. Valid AI approaches would involve fuzzy logic as well as relational databases. Products would include a 3D representation of object (plate map) and its dynamical properties (spin, rotation), bulk (average) as well as maps (spatial distribution) of compositional (mineralogical and elemental) components, interior (gravity and magnetic anomalies) and surface (color, albedo, morphology) properties. Mapping is contingent upon creation of the plate map, primarily from the imager input. Thus, the scheduling, generating and deriving useful data from the imager would be given top priority in terms of resource allocation.

4. Interpreting Data: Human agents interact with Leader to interpret analyses in terms of systematic trends in relationships between properties and produce higher level products in keeping with overall mission goals to characterize a group of bodies and their origin in context of the solar system. Because nothing resembling real-time interaction will be possible at this distance, human involvement will be limited. The leader must be designed to do all of this work without human agents, although human agent involvement would be valuable. This work is deductive, highly intuitive, and model-based. Fuzzy Logic and relational databases would be valuable here as well.

The Impact of Environmental Support and Cognitive Style on Space Explorers

Explorer communities have direct implications for design of Intelligent Software Agents. Of particular significance is high level of strategic autonomy demanded here, as a result of the relatively infrequent contact with human agents. The system must be designed to take full advantage of communication between intelligent agents operating in extremely varied cognitive frameworks whenever it does occur. On the Myers Briggs axis for cognitive validation, end-members range from sensing (S), involving data acquisition, an inductive activity, to intuitive (N), involving use of an a priori model a deductive activity. The perceived goals include not only the more usual avoidance of risks (survival of a significant portion of the explorers is important!), but most importantly, acquisition of valid targets. On the Myers Briggs axis for cognitive activity, end-members range from completion-oriented or judging (J) to process-oriented or perceiving (P). The word valid is very important in this context. Not just any target selected at random will be acceptable. A high level of intelligence is required, involving use of a priori and evolving models as well as real-time data, in a deliberate selection process. Again, some combination of real-time data and evolving model of the population of targets will be used in this assessment as well. Thus, the Decision Support System (DSS) could involve empirical data (DB), current models (MB), or some combination of the two (DMD or MDM). Decision Support System and the Cognitive style of multi-agents synergistically interact to influence the ability to perform essential tasks during exploration, and that these two parameters can be optimized to achieve the goal of valid target acquisition.

Reference available on request from first author. ■

Shawnee Vickery, Feature Editor
vickery@pilot.msu.edu