

CHAPTER 3

SPACE EXPLORATION: THE INTERSTELLAR GOAL AND TITAN DEMONSTRATION

3.1 Introduction

The small Pioneer 10 spacecraft, launched from Earth on March 2, 1972, represents mankind's first physical extension into interstellar space. Having traversed the Asteroid Belt and given scientists their first good look at Jupiter and its satellites, the vehicle now rushes toward the edge of the Solar System at a speed of about 3 AU/yr. The exact moment of penetration into extrasolar space is unpredictable because the boundary of our System is not precisely known, and because the spacecraft's ability to transmit useful data will likely degrade by the time of passage (circa 1986) that it will be unable to report transit of the heliosphere when this occurs.

Several other unmanned vehicles will also eventually exit the Solar System. However, as Pioneer 10 none of these were designed specifically as interstellar probes, and comparatively little work has yet been accomplished with the aim of developing such craft. Still less effort has been directed toward the ultimate goal of manned interstellar exploration.

3.1.1 Automated Interstellar Space Exploration

The most extensive study of interstellar space exploration to date has been Project Daedalus, an analysis conducted by a team of 13 people working in their spare time under the auspices of the British Interplanetary Society from 1973 to 1978 (Martin, 1978). The focus was a feasibility study of a simple interstellar mission using only present technology and reasonable extrapolation of foreseeable near-future capabilities.

The proposed Daedalus starship structure, communications systems, and much of the payload were designed entirely within today's capabilities. Other components, including the machine intelligence controller and adaptive repair systems, require a technology which Project members expected would become available within the next several decades. For example, the propulsion system was designed as a nuclear-powered, pulse-fusion rocket engine burning an exotic deuterium/helium-3 fuel mixture, able to propel the vessel to velocities in excess of 12% of the speed of light. Planetary exploration and nonterrestrial materials utilization were viewed as prerequisites to the Daedalus mission, to acquire useful experience and because

the best source of helium-3 propellant is the atmosphere of the gas giant Jupiter (to be mined using floating balloon "aerostat" extraction facilities). This ambitious interstellar flyby was thought possible by the end of the next century, when a solar-system-wide human culture might be wealthy enough to afford such an undertaking. The target selected for the first flight was Barnard's star, a red dwarf (M5) sun 5.9 light years away in the constellation Ophiuchus.

The central conclusions of the Project Daedalus study may be summarized roughly as follows: (1) Exploration missions to other stars are technologically feasible; (2) a great deal could be learned about the origin, extent, and physics of the Galaxy, as well as the formation and evolution of stellar and planetary systems, by missions of this kind; (3) the necessary prerequisite achievements in interplanetary exploration and the accomplishments of the first interstellar missions would contribute significantly to the search for extraterrestrial intelligence (SETI); (4) a funding commitment over 75-80 years is required, including 20 years for vehicle design, manufacture and checkout, 30 years of flight time, and 6-9 years for transmitting useful information back to Earth; and (5) the prospects for manned interstellar flight are not very promising using current or immediately foreseeable human technology.

A more recent study (Cassenti, 1980) concludes on a more optimistic note: "We are like 19th Century individuals trying to imagine how to get to the Moon. Travel to the stars is extremely difficult and definitely expensive, but we did get to the Moon and we can get to the stars." Cassenti supports the Project Daedalus judgment that only vehicles capable of achieving more than 10% of the speed of light should be examined and that the preferred propulsion system now is "a version of the nuclear pulse rocket for unmanned exploration and combinations of the nuclear pulse rocket and the laser-powered ramjet for propelling manned interstellar vehicles."

Even more imaginative and longer-range interstellar missions of galactic exploration have been considered by Robert A. Freitas Jr., a participant in the present study (Freitas, 1980a, 1980b; Valdes and Freitas, 1980). He concludes that self-reproducing interstellar probes are the preferred method of exploration, even given assumptions of a generation time of about 1000 years and a 10-fold improvement in current human space manufacturing technology. He envisions "active programs lasting about

10,000 years and involving searches of 1,000,000 target stars to distances of about 1000 light years in the Galactic Disk ...” and states that interstellar probes will be superior to beacon signals in the search for extraterrestrial intelligence.

The Space Exploration Team was charged with defining a challenging mission for the next century which could be a technology driver in the development of machine intelligence and robotics. Interstellar exploration was early identified as the ultimate goal, where this would focus on an investigation of planetary systems in the solar neighborhood discovered through SETI operations or by searches with large apodized visual telescopes (Black, 1980) in Earth orbit. Though previous studies of interstellar exploration missions are few, even these clearly suggest the need for high levels of automation.

The Team defined a general concept of space exploration centered on the notion of an autonomous extrasolar exploratory machine system. This system incorporates advanced machine intelligence and robotics techniques and combines the heretofore separate and manpower-intensive phases of reconnaissance, exploration, and intensive study into a single, integrated mission. Such an automatic scientific investigation system should be useful in the exploration of distant bodies in the Solar System, such as Jupiter and its satellites; Saturn and its rings; Uranus, Neptune, Pluto and their moons; and perhaps comets and asteroids as well. It may provide tremendous economies in time, manpower, and resources. Interstellar exploration seems virtually impossible without this system, which is itself a magnificent technology driver because the level of machine intelligence required far outstrips the state of the art (see section 3.3).

This report cannot review the entire gamut of reasons for human interest in the physical exploration of the Solar System and the Universe. Recent space research programs have stimulated large numbers of people from various scientific disciplines to join in the challenge of interplanetary exploration. Astronomers and geologists have participated since they represent the sciences traditionally most involved in the observation and classification of planetological and celestial phenomena. During the last two decades researchers from other physical sciences and the biological sciences have become interested in investigating how the laws of nature operate in the cosmos, using the techniques of radio astronomy and space exploration including direct biological samplings of other planets. Interest in the outer Solar System and deep space will likely remain high among natural scientists.

It is assumed that these reasons, coupled with the seemingly basic need of human beings to satisfy their inherent curiosity when confronted by new environments, are sufficient to motivate the economical exploration programs that advanced machine intelligence systems will make possible. Appendix 3A includes a summary of the ideas of the team’s

student member, Timothy Seaman, whose feelings may be representative of those of the generation of young Americans most likely to receive the first major benefits from mankind’s more ambitious future ventures into space.

Although interstellar exploration was identified as the ultimate goal, detailed mission analyses are not provided. The determination of technological, economic and political feasibility for such complex, expensive, and extraordinarily long-duration undertakings must wait until advanced machine-intelligence capabilities of the type required for an extrasolar voyage have been successfully demonstrated in planetary missions conducted entirely within the Solar System. Accordingly, the major emphasis of the present study is a Titan Demonstration Mission (fig. 3.1) conceptualized to require the evolution of equipment and machine intelligence capabilities which subsequently may be applied to autonomous interstellar operations.

3.1.2 The Titan Demonstration Mission

The demonstration mission concept leads ultimately to development of a deep space - system incorporating advanced machine intelligence technology capable of condensing NASA’s current three investigatory phases -- reconnaissance, exploration, and intensive study -- into a single, integrated, autonomous exploratory system. This should yield significant economies in time and resources over present methods (table 3.1).

TABLE 3.1.— SPACE EXPLORATION: THE INTER-
STELLAR GOAL AND TITAN DEMONSTRATION

<p>Goal: Evolution of capability for autonomous investigation of unknown domain.</p> <p>Approach: Integrate previously separate investigation steps into single mission.</p> <ul style="list-style-type: none"> ● Advanced propulsion capability ● Global scale investigation by remote sensing <ul style="list-style-type: none"> ● Advanced sensors ● Machine intelligence for information extraction and plan follow-up ● Limited number of in situ exploration vehicles ● Autonomous hypothesis formation to classify information and develop new theories

The Space Exploration Team proposes a general-purpose robot explorer craft that could be sent to Titan, largest of Saturn’s moons, as a technology demonstration experiment and major planetary mission able to utilize the knowledge and experience gained from previous NASA efforts. Titan was chosen in part because it lies far enough from Earth to preclude direct intensive study of the planet from terrestrial




Figure 3.1. – Titan Demonstration Mission.

observation facilities or easy teleoperator control, yet is near enough for system monitoring and human intervention as part of a developmental process in the demonstration of a fully autonomous exploration technology. Such capability must include independent operation from launch in Low Earth Orbit (LEO); spiral Earth escape; navigation; propulsion system control; interplanetary flight to Saturn followed by rendezvous with Titan; orbit establishment; deployment of components for investigation and communication; lander site determinations; and subsequent monitoring and control of atmospheric and surface exploration and intensive study. The target launch date for the Titan Demonstration Mission was taken as 2000 AD with 5 years on-site. Knowledge gained from the Titan exercise could then be applied to the design of follow-on exploration missions to other planetary systems.

A number of specific criteria were decisive in the selection of Titan as a premier demonstration site for the autonomous exploration system concept:

(1) Titan is one of the few bodies in the Solar System where the physical and atmospheric conditions are partially unknown and interesting, but also still lie within acceptable tolerance ranges for equipment survivability.

(2) Titan, 9.54 AU distant from the Sun, is far enough from Earth to preclude intensive study using terrestrially based, scientific, experimental, and observational equipment, to deny easy teleoperator operations, and to require fully autonomous systems functioning while still being close enough for monitoring and intervention by humans as the demonstration experiment evolves.

(3) The existence of a heavy atmosphere provides a good test for system flexibility since atmospheric modeling

is crucial in understanding surface conditions and evaluating the possibility of life. Thus, smart multispectral correlation systems development is essential.

(4) The shrouded surface provides an unknown environment in which to test imaging systems without bias.

(5) Titan is better capable of capturing and holding the public interest than other bodies for some of the same reasons that it has received increasing scientific attention; for instance, the fact that it holds a faint hope for lifeforms (past, present, or future) and requires the full NASA array of equipment including the manned Shuttle. The Saturnian moon already has been popularized by Carl Sagan in his PBS television series "Cosmos" with a visually striking simulated Saturn ring penetration and Titan landing, and

Voyager I vastly increased our scientific knowledge of Titan during its encounter with the planet in November 1980.

(6) Precursor missions will provide enough knowledge of Titan and the Saturn environment to allow verification by Earth-based scientists of the atmospheric and surface models sent back by hypothesis-formation modules operating aboard the Titan spacecraft.

(7) A partial knowledge of the Titan environment permits equipment and experiment economies over later missions wherein many more contingencies and hypotheses must be anticipated.

A Titan Demonstration Mission in the year 2000 AD would benefit from two types of heritage (fig. 3.2). The first, knowledge heritage, allows the use of spacecraft com-

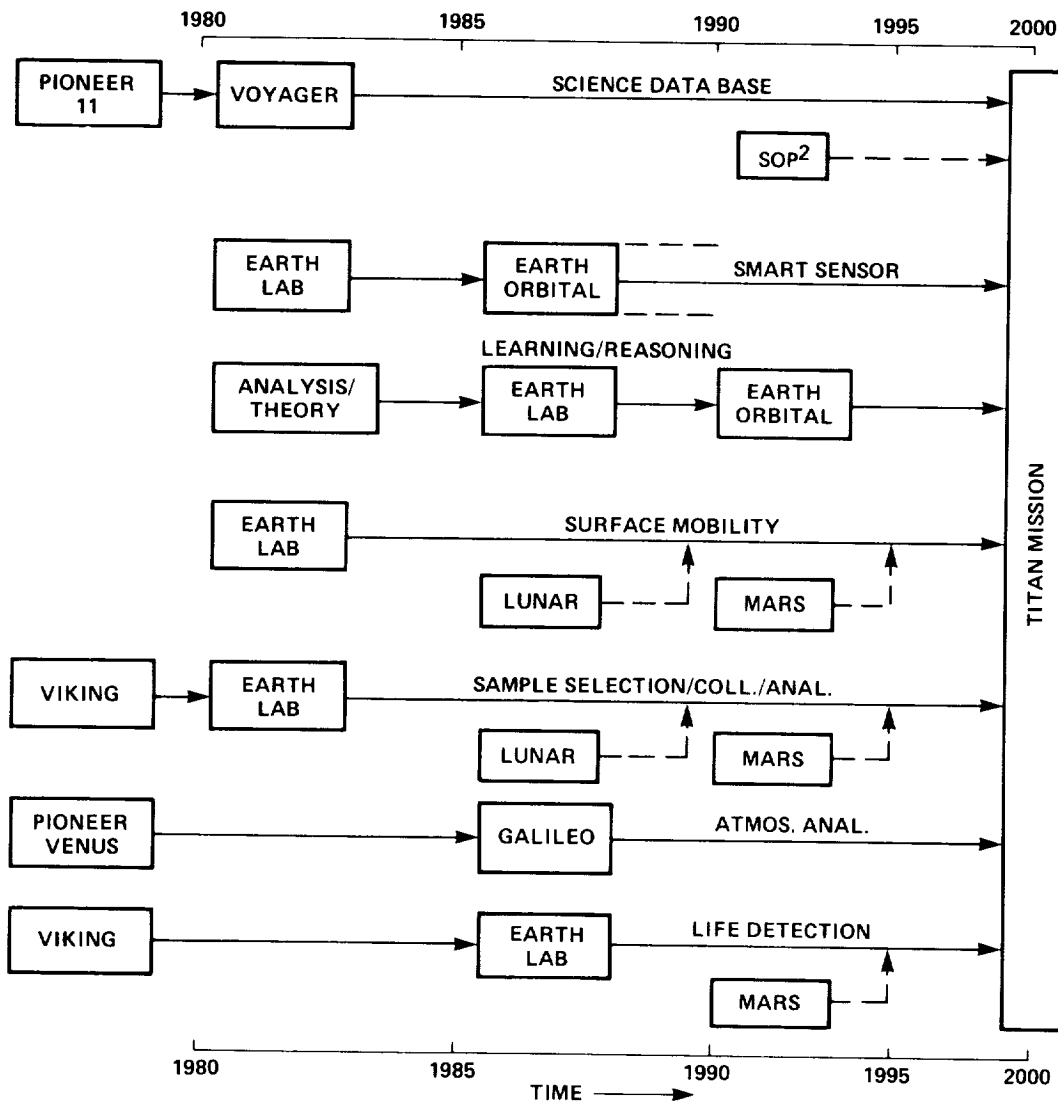


Figure 3.2.— Prior mission contributions to desired Titan mission capabilities.

ponents which need not be designed to cope with wholly unknown alien environments. The experience gained during the Pioneer 11 and Voyager encounters with Saturn and its moons has provided essential prior scientific and engineering data on Titan and its surroundings. The second, equipment heritage, permits investigative techniques developed for earlier missions to be adapted in modified form for the Demonstration. Many pre-Titan spacecraft operations address the same basic objectives in planetary exploration and provide a useful remote-sensing technology base to carry them out. For example, the Viking, Pioneer Venus, and Galileo missions furnish techniques for in situ atmospheric analysis, and valuable experience with surface analyses searching for microbial life was gained during the Viking mission to Mars.

A number of planned or opportunity missions currently under consideration by NASA offer further possibilities for technology development in directions useful for the

Titan Demonstration – such as the proposed lunar and Mars missions employing autonomous surface roving vehicles and advanced methods for sample selection, collection and analysis, and the VOIR (Venus Orbiting Imaging Radar) system for the development of a planetary radar mapping capability. Since the global characteristics of Titan are included within the scope of the Demonstration, opportunities for knowledge and for equipment heritage exist with respect to the proposed Saturn Orbiter Dual Probe (SOP²) spacecraft.

In summary, the proposed Titan technology demonstration experiment and major space exploration mission utilizes the knowledge and experience gained in previous NASA operations. In turn, the Demonstration itself serves as the verifying mission for an autonomous space exploration capability which is the ultimate goal.

Figure 3.3 shows the relationships between the research areas of the four Study Teams and the Titan and interstellar

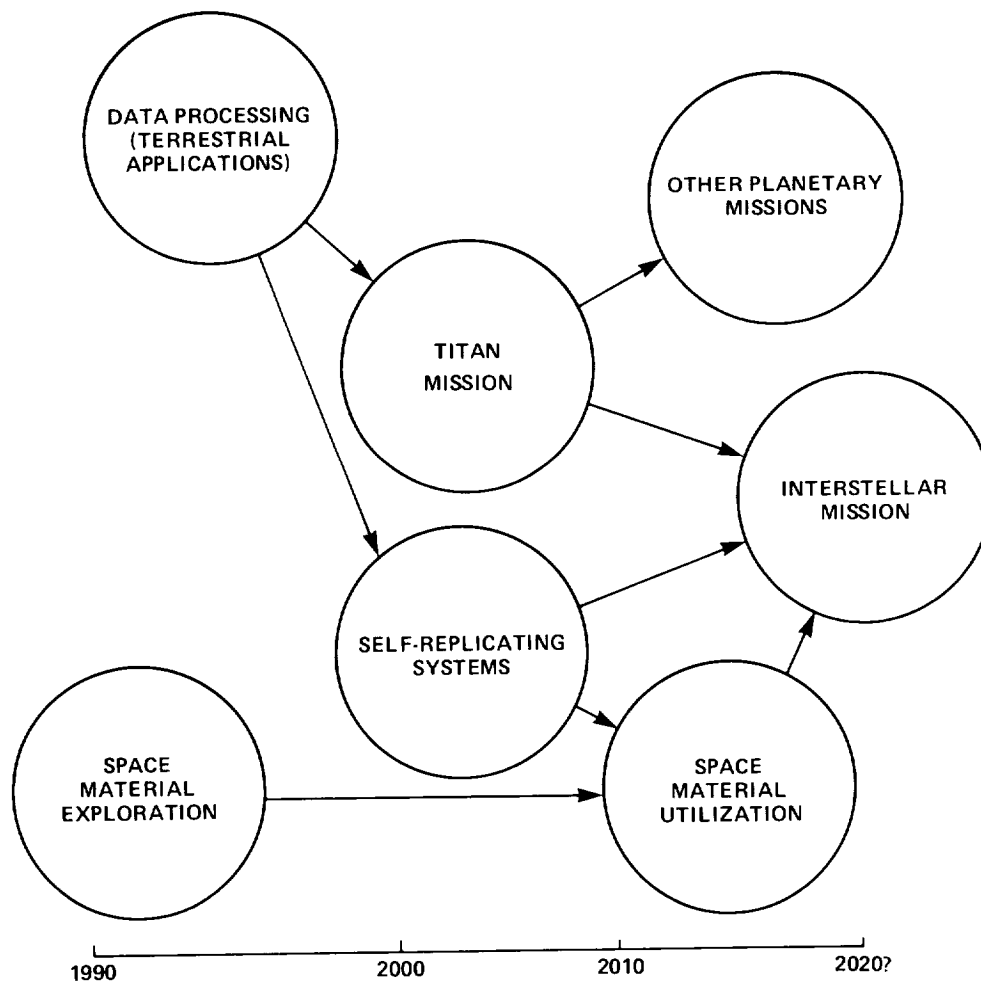


Figure 3.3. – Relationships between space exploration and other 1980 study areas.

mission concepts addressed in this chapter. Of particular interest is the question, "How soon after the Titan mission will extraterrestrial materials be utilized to facilitate interstellar exploration missions?" A Delphi poll was conducted using all Study participants (considered the best sample of experts immediately available to consider the question) and the results were: Median year 2028 AD, with the 14 estimates ranging from 1995 AD through 2100 AD.

3.2 Titan Demonstration Mission Definition

The Titan Demonstration Mission as envisaged by the Space Exploration Team encompasses a continuum of scientific investigative activities culminating in a fully autonomous extrasolar exploratory capability. The primary focus is on condensing into a single extended mission NASA's present sequential approach of reconnaissance, exploration and intensive study. In the past, interplanetary discovery has required Earth-launch of consecutive exploratory devices designed on the basis of data gathered by precursor craft. This approach assumes a broad range of sophisticated sensing equipment but little capability for onboard processing. Analysis of acquired data typically has been relegated to earthbound scientists who make judgments to determine the best next course of action, a procedure which incurs considerable time delays in return transmission of data as well as in ground-based control of distant spacecraft. An even more dramatic delay problem emerges with respect to the deployment of subsequent exploratory devices. In the case of Mars, for example, an initial reconnaissance vehicle (Mariner 4) was dispatched in 1964 but it was more than 10 years later (in 1975) before Viking 1 could be launched to attempt a Martian landing and a more intensive planetary investigation.

Mars, of course, is one of Earth's closest neighbors. Time delays in data transmission and control functions reach a maximum of 21 min in each direction, and travel time from Earth to Mars is approximately 1 year. In the outer Solar System the delay for one-way data transmission and control is measured in hours or days, while at interstellar distances, delay is measured in years with travel times of decades or more. As exploration goals are extended into the farthest reaches of space, development of nontraditional techniques and systems requiring a lesser dependency on Earth-based operations and possessing far greater autonomy become increasingly desirable and necessary. It is in this spirit that the Titan Demonstration Mission is proposed – anticipation of the potential for advanced machine intelligence eventually to permit fully autonomous exploration of the interstellar domain, a capability born of earlier demonstrations within the closer context of the Solar System.

In order to maintain linkages with current and future NASA activities (e.g., Voyager, Saturn Orbiter Dual Probe) and between short- and long-term objectives, the initial Titan demonstration relies upon extensions of current arti-

ficial intelligence (AI) techniques where these are appropriate. For example, by the year 2000 a considerable amount of information about Titan's characteristics, including a basic atmospheric model, already may have been compiled. Assuming research and development progresses in both interacting simulation models and rule-based automated decisionmaking, then extensions of current AI knowledge-based systems will have the potential to contribute to the automatic maintenance of mission integrity to insure the survival of mission functions and components.

To the extent that new developments in machine intelligence technology move in the appropriate directions, the Titan mission might include demonstrations of autonomous onboard processing of mechanically acquired data in at least one sample of scientific investigation. This results in great compression of return information because only the "important" or "interesting" hypotheses about the target planet are transmitted back to Earth. Such a function presupposes a machine capacity both for hypothesis formation and for learning, neither of which is inherent in state-of-the-art AI technology (see section 3.3). Significant new research in machine intelligence is a clear prerequisite to successful completion of the proposed Titan Demonstration Mission (see table 3.2).

TABLE 3.2.— TITAN EXPLORATION MISSION DRIVERS

Technology
<ul style="list-style-type: none"> – A coordinated surrogate scientific community on and around Titan – Long system life – 10 years or more reliable/redundant propulsion/energy – Distributed decision and expert systems – Self-monitor and repair ability – Semi-autonomous subsystems <ul style="list-style-type: none"> Probes, Landers, Rovers, Satellites – Data storage and reduction; information communication to Earth – Integrated multisensor capability
Intelligence
<ul style="list-style-type: none"> – Overcome the intelligence barrier. Current AI capabilities and research will not achieve autonomous MI needs for space exploration – MI for space exploration must be able to learn from and adapt to environment. To be able to formulate and verify hypotheses is essential, but may not be sufficient.
<p>Goal: Full autonomic exploration system with human intervention option.</p>

While Titan is too distant to explore efficiently using traditional methods it is still near enough to monitor the performance of automated functions and to take intervening action should the need arise. As exploration distances extend farther out into the Solar System, such intervention becomes increasingly difficult so the demand for greater mission autonomy and higher-level machine intelligence rapidly intensifies. An outline of operational mission stages integral to the full range of exploratory activity, from the Titan demonstration to interstellar exploration, is presented below. Each phase underscores a variety of machine capabilities, some unique and some overlapping, required if full autonomy is to be achieved. These capabilities represent the primary technology drivers for machine intelligence in future space exploration.

3.2.1 Titan Mission Operational Stages

A fully automated mission to Titan (and beyond) requires a very advanced machine intelligence as well as a system which is highly adaptive in its interactions with its surroundings. This latter aspect is even more significant in extrasolar missions because a sufficient operational knowledge base might not be available prior to an encounter with new planetary environments. The explorer must generate and use its own information regarding initially unspecified terrain, and this knowledge must evolve through the updating of databases and by the continual construction and revision of models. Such a machine system should be capable of considerably higher-order intelligent activities than can be implemented with state-of-the-art techniques in artificial intelligence and robotics.

The short-term mission objective is to encompass the tripartite staging of NASA missions within a single, fully automatic system capable of performing scientific investigation and analysis, the immediate objective being a complete and methodical account of Titan. Later, and as a longer-term goal, given the successful achievement of the short-term objective, a similar exploration of the outermost planets and bodies of the Solar System could be conducted with improved equipment, building on the systems operations knowledge gained at Titan.

The proposed exploration system must be capable of the following basic functions:

- (1) Select interesting problems and sites.
- (2) Plan and sequence mission stages, including deployment strategies for landers and probes.
- (3) Navigate in space and on the ground by planning trajectories and categorizing regions of traversability.
- (4) Autonomously maintain precision pointing, thermal control, and communications links.
- (5) Budget the energy requirements of onboard instrumentation.
- (6) Diagnose malfunctions, correct detected faults, and service and maintain all systems.

(7) Determine data-taking tasks, set priorities, and sequence and coordinate sensor tasks.

(8) Control sensor deployment at all times.

(9) Handle and analyze all physical samples.

(10) Selectively organize and reduce data, correlate results from different sensors, and extract useful information.

(11) Generate and test scientific and operational hypotheses.

(12) Use, and possibly generate, criteria for discarding or adopting hypotheses with confidence.

One way to formalize the precise characteristics of a proposed mission is in terms of a series of prerequisite steps or stages which, in aggregate, capture the nature of the mission as a whole. The operational mission stages selected for the Titan demonstration analysis are: configuration, launch, interplanetary flight, search, encounter, orbit, site selection, descent, surface, and build. Each is discussed briefly below.

Configuration. This initial phase addresses considerations of size, weight, instrument specifications and other launch vehicle parameters, and usually depends on the equipment and tasks required for a specific mission. Questions concerning the precise nature of the investigation and experimentation traditionally are taken up at this point.

For deep-space exploration, spacecraft configurations must be general and flexible enough to handle a wide range of environments. Hardware and software impervious to extreme pressure, temperature, and chemical conditions and with long lifespans are required. Also, a diverse assortment of onboard sensors with broad capabilities is necessary to produce basic information via complementary and selective sensing to be used in scientific investigation and planning.

Launch. The focus of this stage depends to some extent on the perceived configuration of the mission vehicle. Issues related to propulsion and energy needs and appropriate launch sites (e.g., Low Earth Orbit vs vicinity of extraterrestrial resources utilized for the mission) are decided. The launch phase is conducted largely by Earth-based humans, but could benefit from machine intelligence capabilities (e.g., CAD/CAM/CAT) for testing, checkout, flight preparation, and launch support.

Interplanetary flight. Prior to Viking and Voyager, unmanned flyby and orbiter spacecraft were totally dependent upon Earth-based remote observation and direct human intervention to accomplish accurate navigation, stationkeeping, and rendezvous and docking maneuvers (Schappell, 1979). This underscores the control and communication time delay problem that limits efficient investigation of distant bodies such as Titan and even more dramatically constrains exploration of the interstellar realm. Some ground-based support for the initial Titan

Demonstration Mission may be appropriate in computing navigational corrections, but subsequent deep-space exploration requires a fully autonomous navigation system. Such systems also improve cost-effectiveness by reducing the amount of ground support necessary to accomplish the missions. Potential savings in equipment complexity, operational costs, and processing time will motivate the development of autonomous systems for near-Earth and deep-space vehicles.

Consider, for instance, the Viking mission, one of the most complex interplanetary operations attempted to date. The Mars landers were remotely operated robot laboratories equipped with comparatively highly automated instrumentation. Many spacecraft functions could be performed adaptively, accommodating to changing necessities during the mission. Even so, the operational system required major navigational changes to be specified 16 days before the indicated flight action. Several hundred people on Earth were involved in science data analysis, mission planning, spacecraft monitoring, data archiving, data distribution, command-sequence generation, and system simulation. An infusion of advanced machine intelligence could significantly reduce this major mission cost.

In addition to navigation, the spacecraft also must maintain attitude and configuration control, thermal control, and communications links. These functions involve the use of feedback loops and built-in test routines. One way to visualize a greatly improved system is to conceptualize a machine intelligence capable of sequentially modifying its activity as a result of experience in the environment, with an additional capability of internalizing or "learning" the relationship between environmental states and corrections to guide future modifications and coordinate them with anticipated states. Such goal-directed intelligent functioning is not possible with state-of-the-art AI technology. However, it is conceivable that a machine system could be provided with a capacity to represent its present state, some goal-state of equilibrium or stability and a means of noting and measuring any discrepancy between the two, and, finally, effectors or actuators for modifying the present state in accordance with the programmed goals.

Search. During the Search phase the system performs preliminary analyses while approaching the target body. The information acquired is integral in making decisions about subsequent activities as well as the point at which to begin preliminary analysis. The spacecraft must be able to employ appropriate sensing equipment to collect raw data and to modify sensor utilization as a result of feedback information. Inherent in this formulation is the capacity of the system to perform some analysis using the raw data it has collected and to make decisions about mission sequencing based on analysis results.

Complementary and concurrent sensing tasks are scheduled according to the time required for their completion,

the point at which their output becomes important to ongoing model construction, and the relative importance of the results. Another significant factor is spacecraft-instrumentation power scheduling, assuming that the supply of energy is insufficient to allow all subsystems to operate simultaneously. Scientific tasks must be scheduled to take into account possible mission-control functions that might override them. Collection tasks producing data having multiple uses or particular utility in mission integrity operations (self-maintenance, survival, and optimization) have high priority. All operations are to be accomplished without benefit of direct human intervention.

For the initial Titan mission, one might attempt to automate all search functions by means of an onboard expert system that utilizes known information about the conditions on Titan and that is capable of examining and choosing from among preselected resident hypotheses (leading finally to some judgment as to what action to take based on probability calculations). However, such a system could be highly fallible because information gaps and inaccuracies in its available range of hypotheses might lead to serious misjudgments. In the case of the long-term objective — interstellar navigation — the consequences of an incomplete knowledge base are even more dramatic. The team concludes that expert systems of the current AI variety cannot satisfactorily perform the Search task.

One possible solution, and a potentially valuable technology driver, is an advanced type of expert system able to update and modify its own knowledge base as a result of experience — that is, as a result of the analytical actions which it performs on its own environment. On Earth the advent of such an advanced system would eliminate time-consuming and costly human analysis and reprogramming typical of state-of-the-art expert systems (which would be particularly inefficient in space applications where huge time delays often must be accommodated). Self-modification of advanced expert systems also prepares the exploration system to make autonomous decisions and corrections regarding its relationship with the environment.

An additional essential task *en route* to an unknown planetary system around another star is the determination of gross parameters such as sizes, masses, densities, orbital periods, rotational periods, axial tilts, and solar distances for each member planet and moon. A fully autonomous spacecraft would utilize these characteristics, determined by early data collection, in making onboard selections of appropriate bodies to explore.

Given the existence of specific atmospheric conditions determined by long-range remote sensing, logical hypotheses may be generated to predict the surface conditions of the chosen celestial body in terms of the possibility of life and the compatibility of the planet with spacecraft hardware and engineering. Decisions must then be made on the basis of preliminary analyses whether to proceed and establish orbit around the planet for further exploration, or

to choose another target. An intriguing alternative would be a system capable of redesigning or adapting its equipment to accommodate the relevant alien environmental conditions.

Encounter. The processing of image data is probably one of the most computationally demanding tasks performed during planetary exploration missions. In the Encounter phase, when the spacecraft controller must make a quick go/no-go decision on the question of orbital insertion, the data processing challenge includes speed as well as volume. The problems of distance and communications delays, coupled with the necessity of making rapid local decisions, virtually demand that image analysis during Encounter be accomplished by fully autonomous onboard processing systems.

One possibility is an imaging system capable of describing a planetary body much as an astronaut would. For example: "The surface is bluish with some brownish areas near the equator. There appear to be thin wispy clouds covering a 100×200 km area centered about 75° N and 30° W." The observation of "bluish" and "brownish" indicates the processor's ability to match raw data inputs to color concepts understood by humans. The identification of "wispy clouds" suggests the capability of matching data in a sequential region of the image to the known concept of "wispy." The ability to match regions, spectral data, and other features in an image to stored concepts in memory requires a reasonably high level of machine intelligence.

Another part of the description of the image observed by spacecraft sensors locates the "wispy" area at a given latitude and longitude. To do this, the processor must be able to establish the geometrical shape of the body encountered and to apply a coordinate system to it. Once this coordinate system is computed it forms the cartographic grid to which all surface features are mapped. While this is a well-understood mathematical procedure, the "number crunching" load is significant and must be executed very rapidly during the Encounter phase.

Orbit. When preliminary analysis suggests a reasonably benign environment warranting further investigation, orbit is established to conduct a more detailed study. The establishment and maintenance of orbital position, like most of the functions already mentioned, should be a fully autonomous process with characteristics similar to the autonomous interplanetary flight navigation system. Onboard automated decisionmakers determine an optimal orbit using information gathered during preliminary analyses, and orbital insertion is achieved.

Multisensor analysis is implemented concurrently with the establishment of orbital position, permitting a more comprehensive investigation of planetary characteristics than during Encounter. During Orbit phase a variety of

sensors and sophisticated image processing techniques are employed to examine atmospheric and surface conditions. Analyses should be conducted both in the context of (1) pragmatic decisionmaking, including assessments of atmospheric pressure, density, and identifications of surface conditions to be utilized in judging which equipment to deploy, and of (2) scientific investigation, such as information acquisition for hypothesis generation.

For the Titan mission an advanced expert system may be used to form judgments about appropriate exploratory equipment for specific environmental conditions. For instance, when deploying probes or landers smart sensors might first assimilate data regarding atmospheric density and pressure. The advanced expert system could then make probability judgments as to how fast probes should fall and how much retrorocket energy is required for landing. Additional assessments could be made of surface conditions, such as whether the surface is composed of a solid, liquid, or gaseous base. This information supports subsequent decisions about necessary configurational requirements of landing craft (e.g., should it be a wheeled, walking, hovering, or floating vehicle?). The above machine intelligence applications could probably be developed on a relatively short-term basis, utilizing minimal extensions of state-of-the-art AI techniques.

In the deployment of such exploratory mechanisms as atmospheric and surface probes, balloons, and landers, intelligent coordination of autonomous orbit maintenance and control is crucial. Since deployment of onboard equipment alters the total mass and mass distribution of the orbiter, some simultaneous revision of the altitude control function, ideally based on "anticipatory information," is required. That is, the spacecraft must anticipate changes in its state prior to component deployment and be prepared to adapt to concomitant variations in its physical state (a specific example of the type of feedback system required to maintain mission integrity).

A much more serious problem for development in the area of machine intelligence is the scientific analysis of data and the autonomous formulation of hypotheses and theories. Current expert systems technology cannot generate and test unique hypotheses that have not been preprogrammed by a human operator. This limitation restricts an exploratory device based on state-of-the-art AI to data analysis, categorization, and classification in terms of existing structures of thought or taxonomies of knowledge. However, in alien environments, particularly those accessible in an interstellar mission, pre-formed scientific notions may not reasonably be applicable; on the contrary, they may serve only to distort higher-order understanding of incoming data. Thus, a major technology driver is the development of an advanced machine intelligence system capable of reorganizing rejected hypotheses, integrating that information with data acquired through sensory apparatus, generating new hypotheses which coordinate all

existing information, and, finally, testing these hypotheses in some systematic fashion. (See appendix 3B for a hypothetical illustration of this point.)

Lander site selection. During this phase some form of mobile surface device compatible with local environmental conditions is deployed according to planetary orbiter directives. This device performs in situ surface and geologic data acquisition, imaging, and representative physical sample collections. Its deployment requires the selection of appropriate landing sites, a major task for the autonomous exploration system controller.

Processed image data of planetary surface conditions permits a mapping of topographic surface characteristics with respect to terrain configuration – a cataloguing of mountains, craters, canyons, seas, rivers, and other features to be correlated with maps of temperature, moisture, cloud cover, and related observables. These maps become the basis for a determination of optimal landing locations. Site selection analysis also must include some judgments regarding areas of greatest “interest” for investigation, necessitating some means of detecting regions of the environment which are anomalous with respect to expectations based on prior preliminary analyses of the locale. Criteria for site selection, as for example geological significance or the possibility of lifeforms, are stored in memory. Imagery to be compared to this set of criteria could be obtained from a world model (see chapter 2) developed during the orbital phase, an application ripe for machine intelligence technology development.

Hazard avoidance at the landing site and terrain traversability for mobile landers are additional considerations in the site-selection process. Some mechanism for self-preservation should be included so that an assessment of potential landing sites is made according to whether they pose a danger or are benign. Only then can adaptive action patterns be undertaken with some reasonable expectation of success.

Descent to surface. The descent to surface should be fully automated even in relatively near-future explorations of the Solar System. Autonomous feature-guided landing poses a unique challenge to image-processing technology. For instance, during a parachute descent the target landing site must be located and tracked by an image processor. As the assigned target is tracked, the lander parachute must be manipulated to steer toward the target much like a sports parachute. While the tracking task is not conceptually difficult, the processing speeds required do not exist in present-day computer hardware. As the surface draws closer, the potential landing site must be reexamined for obstacles hazardous to the craft. This presupposes some stored knowledge of precisely what could pose a hazard, as well as the ability to act upon that information. In the Descent phase, machine intelligence integral to the surface exploration system will require high-accuracy processing and ultra-high speed hardware.

On the surface. Once surface contact is achieved the most interesting and probably the most difficult image processing begins. Self-inspection for damage comes first, followed by verification of the lander’s position. This may involve comparing the surrounding scene with possible projected scenes assembled from the world model, or the analysis could be based on tracking by the main orbiting spacecraft. Next is the planning, scheduling, and commencement of experiments. All conflicts must be comprehended and resolved. If one experiment calls for rock density measurements and no rocks are within reach of the lander’s end-effectors, a decision must be made to schedule another experiment or to move the lander. Such operational decisions require intelligent scene analysis and concept/theory matching.

If preliminary analyses suggest that further investigation is warranted and safe, the lander system for image processing of the surrounding area is deployed. This accompanies the collection of local temperatures, pressures, and general ambient conditions data, as well as sample collection and analysis. To provide these functions the lander (an intelligent robotic device) is equipped with a wide variety of sensor and end-effector apparatus. Vision is especially important for obstacle avoidance and mobility. Stereo vision may prove an invaluable aid in successfully traversing three-dimensional spaces, and also an important safety feature for avoiding depth hazards.

Mobile lander data collection responsibilities require several specific machine intelligence capabilities including (1) pattern recognition to correlate visual images and to detect similarities and differences among data alternatives and (2) decisionmaking to determine whether a particular datum is worth collecting. While it is conceivable that minimal extensions of state-of-the-art expert systems might prove adequate to address the problem of datum “worth,” still there remains a sizable gap between current capabilities in computer perception (pattern recognition) and capabilities needed for tasks integral to the proposed mission – another crucial technology driver.

While some of the Titan mission performance demands on robot manipulators are not as critical as on industrial assembly lines, still there are definite constraints. Spacecraft effectors must operate in completely unstructured environments unlike state-of-the-art factory robots which move only in small, comparatively well-defined work areas. Precision requirements are fairly modest for explorer manipulators when they are handling physical samples, but placement accuracy must be considerably improved whenever the system is responsible for joining closely fabricated pieces during instrument repair, component reconfiguration or construction. Manipulator supervision is supported primarily by visual sensing, though a wide variety of other sensor inputs may supplement optical techniques.

A potentially difficult image processing task is the coordination of manipulator movements with those of the

target object, better known as "hand-eye coordination." Image processing may be used accurately to find the position (in three-dimensional space) of the object to be manipulated as well as the grasping surfaces of the manipulator itself. Locating these surfaces might involve matching the received images with memorized models (or concepts) of the object and the end-effector, a tremendous challenge to present-day machine intelligence technology. An alternative method requires using pressure- and force-feedback, as well as proprioceptive information (sensory input designating body or effector orientation) to reduce image processing requirements.

Movement of the lander demands that a safe, obstacle-free path be found across the landscape. This may entail generating a contour map of the surface surrounding the lander (perhaps using high-resolution satellite/orbiter data) and derivation of a clear path from this map. State-of-the-art laser scanning techniques already have proven adequate to handle the task of topographic analysis for purposes of local wild-terrain locomotion. Hazards hidden from view along the intended itinerary must be identified *en route*, and the path ahead continually re-scanned and updated as in the case of a human walking through a rocky area.

An alternative (and more difficult) approach places greater reliance on autonomous lander processing systems. A planet model provides an apparently traversible path from the landing site to another location observable from the landing site (based on low-resolution data). This "fuzzy" trail is given to the lander controller which then must negotiate its own path from the first position to the second, must identify and work its way around such obstacles as gulleys, creeks, or rubble invisible in the low-resolution model. In addition, during each traverse the lander analyzes the surrounding scenery and searches for significant or unusual objects while also keeping track of its location. Thus, a great deal of image processing and map updating must be done that requires formidable onboard computing power, as well as advanced machine intelligence techniques.

Build. The Build phase actually lies in the domains of space manufacturing (chapter 4) and machine replication (chapter 5), but nevertheless, is worth mentioning here as an important prerequisite for extending the proposed mission to intensive Solar System and interstellar exploration. At some (yet undefined) point it becomes necessary to provide machines with mining, materials processing, construction, repair, and perhaps, even replicative capabilities in order to escape the enormous cost of building and launching burgeoning masses of exploration equipment from Earth (Freitas, 1980b). With respect to the Titan Demonstration Mission, a first step toward the ultimate goal of machine self-sufficiency would be an onboard provision for machine hardware components with the

ability to make adaptive modifications to the system as a result of preliminary analyses of probe and landing craft needs.

3.2.2 Scientific Investigation: Remote Sensing and Automated Modeling

The concept of space exploration presented above suggests the potential capability of an interstellar spacecraft to develop complete detailed models of planets and moons in other solar systems and to return these to Earth as major scientific discoveries about the Galaxy. These models would include information about the planets' atmospheres, surfaces, subsurfaces, electromagnetic and gravitational fields, and any evidence of lifeforms.

Having first characterized the operational mission stages and identified the important machine intelligence requirements of each, the Space Exploration Team chose to consider at greater length one aspect of the Titan Demonstration system capacity to conduct useful scientific investigations: automated modeling of an unknown celestial body. This particular aspect of the scientific investigation capability was selected because it involves the full range of high-level machine intelligence required for autonomous space exploration, while simultaneously relating to the orbit-based world model deployment scheme contemplated by the Terrestrial Applications Team (see chapter 2).

In terms of the preceding discussion of the operational phases of space exploration missions, the task of creating such models is the first and foremost task of the Orbit stage. Detailed remote sensing is undertaken in the mission orbital phase to complete atmospheric modeling and to map various physical parameters of the surface. Perhaps as much as 90% of the total information gathered in the exploration of an unknown body can be collected by the orbiter.

A complete world model describes atmospheric and surface physical features and characterizes the processes which govern the dynamic states of the planet and its atmosphere. The job of constructing a world model may be broken down into two separate categories: building an atmospheric model and examining processes in the surface environment, described below. Since a great deal of work is under way at NASA and at various universities in the analysis of Landsat and weather satellite information, it can be anticipated that much of the groundwork in the techniques for assembling a planetary model will have been laid long before deployment of the Titan mission. Not only is the development of a terrestrial world model an essential precursor research program in pursuit of interstellar mission technical requirements, but it also provides valuable Earth resource information in the more immediate future. Creating and automatically modifying world models based on inputs from a variety of sensors is a machine intelligence technology in which research should be encouraged.

Atmospheric modeling. An accurate atmospheric model is essential to successful landing, scientific analysis, and the prediction of the possibility of indigenous life. The construction of an atmospheric model for Earth (including composition, structure and dynamics) has taken many years, an iterative process dictated by evolving technology plus the developing knowledge and expertise of investigators in a young field. To a large extent this emerging methodology has been driven by the measurability of accessible variables, which may or may not be optimal from a systems theoretical point of view. But given higher technology, observational freedom from Earth's atmosphere, and fresh unknown territory to explore, many more options become available with respect to what should be measured and in what order to define an atmosphere most efficiently and unambiguously. The process has not yet been adequately systematized to permit clear-cut rational choices.

Atmospheric modeling should begin early in the approach to an unknown planet since many mode-of-exploration decisions require information on the nature of the atmosphere. During the course of the mission the atmospheric model accumulates greater detail with continuous updating as higher sensor resolution is achieved and probes are deployed for direct measurements. The investigation of an atmosphere differs from studies of surface characteristics in that it involves the complex integration of many inter-related subhypotheses and measurements of numerous allied parameters. Studies of the surface are more a problem of deriving hypotheses from completed maps representing different measurements and then overlaying these maps as a final step.

Specific initial tasks related to atmospheric modeling include:

- Determination of the region of the spectrum in which most of the electromagnetic radiation is emitted.
- Determination of the sources of opacity for selection of optimum communications link frequency (for landers and probes) and for choosing wavelengths in which to perform infrared and millimeter radiometry.
- Search for unbroadened spectral lines above the atmosphere to provide information on the overall composition of the air.
- Observe where spectral lines interfere with blackbody temperature measurements and determine the wavelength(s) at which the atmosphere may be fully penetrated and planetary surface temperatures accurately recorded.
- Perform preliminary temperature and pressure measurements, to be updated once a comprehensive atmospheric model has been constructed.
- Begin atmospheric modeling with remote sensing at millimeter and infrared wavelengths.

Surface modeling. The best method for planetary surface structure hypothesis formation requires scanning of the body with sequentially increasing resolution in at least four distinct steps. The first step obtains global average values for temperature, surface structure, composition, etc., and establishes norms for keying future observations at higher resolution. Gross features such as lunar maria and highlands or the martian polar caps would appear in this type of survey.

The second observational phase exposes finer detail, identifying regions on the scale of the Tharsis Plain of Mars or the Caloris basin of Mercury. As the explorer approaches Titan, higher-resolution observations of the surface become possible and morphological changes can be observed in each succeeding frame. Recognition of features such as craters, mountains, rivers, and canyons may be accomplished by an advanced expert system which includes models of surface processes in its knowledge base, although present-day pattern recognition and vision systems will require significant refinement before this capability can be realized.

The third step is the recognition of sites with high potential for usefulness in the construction of world models. Such sites mainly include unusual features that are interesting because of their anomalous nature. Identification requires a stored concept of "usual," as for instance: "There is usually a sharp boundary between continents and oceans" and "Craters viewed from directly above usually are circular." An original supply of these simple concepts are programmed into the system by humans before the mission begins; however, additional and revised definitions of normality must be developed and refined as the mission study of a particular planetary body progresses, with self-developed concepts of "usualness" updated by the system as various stages and modes of multisensor investigation are completed. The recognition of that which is "unusual" is discussed at greater length below.

The fourth and final step includes detailed surveys at maximum resolution of selected sites and additional imaging of various undistinguished sites spaced along a grid to pick up interesting features missed by other searches at lower resolution.

Automated selection of interesting sites. It is desirable to minimize raw data storage in order to maximize the efficiency of onboard concurrent mission tasks and analyses. Some method must be found to deal with the information overload which might result from exhaustive exploratory surveys, particularly high-resolution topographic mapping.

Data preprocessing and compression are needed not only because of memory limitations but also to help reduce the complexity of information to be assimilated into world models. Without some way of narrowing the field of interest or of identifying "highlights," the task of converting multiple correlations of many detailed data sets into complete models is cumbersome and impractical. Simplification

also is needed to perform initial but fairly exhaustive searches for sites warranting further investigation (e.g., potentially interesting, safe for landing, etc.). This activity requires a high-level machine intelligence system able to make good choices of which high-resolution data to save and which to throw away.

One possible approach to the selection problem is for the first mapping system to earmark data anomalies for surveillance at higher-than-normal resolution in subsequent surveys. Anomalies are sought by making detailed comparisons of successive maps of the same region. Alternatively, scan data can be searched for locales in which the measured parameter deviates significantly from a predetermined norm. In practice, single-pixel measurements might be saved if values exceed specified thresholds. Data also may be saved in map regions where measured parameter gradients are as steep or steeper than the slope defined as a "significant" edge for that type of measurement. Effectiveness using norm comparisons depends upon appropriate thresholding, whereas the detection of anomalies by successive map matching may be a less subjective approach. (Measured quantities in either case might include rock types, textures, slopes, temperatures, gas concentrations, symmetries, and colors.)

Abridged maps of anomalies detected using the initial survey maps are then compared in a search for correlations that might identify interesting sites. Both the most common and the most unusual sites have high priority for further examination. Areas of interest are identified and ranked for intensive investigation according to the total number of different types of "edges" they contain. The degree of uniqueness of any given site is a criterion for prioritizing follow-up studies. The occurrence of more than one site exhibiting the same edges or of locations with similar correlations indicates a need for additional study. The distribution of correlated sites might suggest some common factor among them; for instance, latitude or regular temporal variations.

The above method of data analysis is one way of focusing on a few features or locations most useful in constructing world models. These high-value sites are identified by their discontinuous character as compared to their surroundings. The method should correctly report features that would be included in an eyewitness description of the celestial body given by a human observer. For example, the crisscrossing lines of Europa might be singled out, as well as the canyons and streambeds of Mars.

Such striking features are necessary but hardly sufficient to specify an entire planet. A complete view also must include: (1) Large-scale structure illuminated by lower-resolution mapping (e.g., the overall smooth surface of Europa), (2) the construction of models inferred from surface mapping data (e.g., that the cracks and smooth surface of Europa indicate a young, active crust), and (3) the incorporation of atmospheric modeling. A machine intelligence

system which can quickly single out and characterize important features of celestial bodies is required. This system should first be tested with known bodies to verify its ability to rediscover what humans would consider to be significant.

Two approaches for model formation have already been presented. The first applies to atmospheric determination – specifically, a process of multiple iteration and revision starting with a distant view of the unknown world. The second applies to surface environment studies after the explorer system has entered orbit. Both methods correctly recognize that models are not effectively constructed by trying to answer all questions for every pixel. Hypothesis formation is followed by a process of testing and checks in each approach.

3.2.3 Titan Mission Components Concepts

The comprehensive exploration and intensive study of Titan will require an appropriate system of spacecraft components. In this section, preliminary technical specifications are provided for each candidate spacecraft function involved in the Titan Demonstration Mission. Of course, final system configuration is dependent upon progress in machine intelligence techniques and on advances in hardware technology that may occur. The technical level of the following specifications is compatible (though not presented in the same format) with the NASA Space Systems Technology Model. In each case, criteria of maximum cost-effectiveness and minimum equipment proliferation are applied.

Rather than discuss every detailed hardware requirement, the Space Exploration Team elected to focus primarily on aspects of the proposed mission which demand significant advances in current technology. Consequently, and also because the design is largely conceptual, the following quantitative and qualitative information would be ranked "level 3" (relatively low confidence) in the Space Systems Technology Model.

The general features of the Titan Demonstration Mission are given in table 3.3. Table 3.4 lists the candidate spacecraft system elements, including the typical number of each type, their operational locations once deployed, and

TABLE 3.3. GENERAL FEATURES OF THE TITAN DEMONSTRATION MISSION

Status: Opportunity Mission (not in current NASA plans)
Lifetime: 10 years; includes 5 years at Titan
Launch/transfer vehicle: Shuttle/400 kW Nuclear Electric Propulsion (NEP)
Operational location: Titan, Saturn's largest satellite
Total mass: 13,000-17,000 kg
Total power: About 400 kW

mass and power requirements. The major mission accomplishments expected of each system component are shown in table 3.5.

surface mobility, and physical sample selection, collection, and analysis. Other candidate system elements have more specialized functions, the management of which can be

TABLE 3.4.— CANDIDATE SPACECRAFT FOR THE TITAN DEMONSTRATION MISSION

Spacecraft type	Typical number	Operational location	Mass, kg	Power, kW
Nuclear electric propulsion	1	Earth to Titan orbit	10,000 ^a	400
Main orbiting spacecraft	1	Circular polar Titan orbit at 600 km altitude	1,200	... ^b
Lander/Rover	2	Surface	1,800	1
Subsatellites	~3	One at a Lagrange point; others on 100 km tethers from NEP	300	0.3
Atmospheric probe	~6	Through Titan atmosphere to surface	200	0.1
Powered air vehicle	1	Atmosphere	1,000	10
Emplaced science	~6	Surface	50	0.1

^aDoes not include propellant.

^bUses NEP power.

The minimum duration of Titan operations is 1 year. While this would be barely sufficient to complete a nominal mission, it is a short time in comparison to seasonal changes in the Saturn system. (Saturn's solar orbital period is 29 years.) The most significant seasonal effects may be expected within about 5 years of the solar equinox of Saturn and Titan — which occurs in 1980, 1995, and 2010 AD. Hence, the preferred arrival dates are 2005 or 2010 AD, with a nominal mission duration of 5 years. Adding 5 more years for interplanetary flight, the preferred Earth-launch dates are 2000 or 2005 AD.

The success of the Titan Demonstration Mission depends on two essential elements — (1) the main orbiting spacecraft and (2) the lander/rover — and on the machine intelligence which they possess. High-level AI capabilities are needed by the main orbiter to coordinate other system components and to conduct an ambitious program of scientific investigation, and are required by the lander/rover to complete its tasks including safe and accurate landing,

assumed, at least in part, by advanced sensors and machine intelligence aboard the orbiter or landing craft.

Nuclear electric propulsion. The early phases of the mission, beginning with launch from Earth and continuing through Saturn arrival, require a high-performance propulsion system which can deliver the payload within a reasonable flight time (4 to 6 years). Low-thrust Nuclear Electric Propulsion (NEP) is the preferred technology for this purpose. The entire NEP system can be delivered to LEO, then be used for spiral escape from Earth, Earth-to-Saturn transfer, for Titan-rendezvous from a circular orbit around Saturn, and finally for spiral capture into Titan orbit and all subsequent spacecraft orbital adjustments. The main orbiter spacecraft and the NEP system share responsibilities for navigation, guidance, control and sequencing, system monitoring, and communications with Earth.

NEP technology has been studied for a long time but has no current planned application beyond possible cargo transport operations from LEO to Geosynchronous Earth Orbit (GEO). Planetary missions such as the proposed Titan

TABLE 3.5.—TITAN MISSION SPACECRAFT ACCOMPLISHMENTS

Spacecraft type	Possible accomplishments
Nuclear electric propulsion	Spiral escape from low Earth orbit; interplanetary transfer to Saturn; rendezvous with Titan; and spiral capture into 600 km circular polar orbit.
Main orbiting spacecraft	Automated mission operations during interplanetary and Titan phases: this includes interfacing with one supporting other spacecraft before deployments; deploying other spacecraft; communicating with other spacecraft and with Earth; studying Titan's atmosphere and surface using remote sensing techniques at both global characterization and intensive study levels; and selecting landing sites.
Lander/Rover	Lands at preselected site, avoids hazards; intensive study of Titan's surface; selects, collects and analyzes samples for composition, life, etc., explores several geologic regions.
Subsatellite	Lagrange point satellite monitors environment near Titan and is continuous communications relay; tethered satellite measure magnetosphere and upper atmosphere properties.
Atmospheric probe	Determines surface engineering properties and atmospheric structure at several locations/times.
Powered air vehicle	Intensive study of Titan's atmosphere; aerial surveys of surface; transport of surface samples or surface systems.
Emplaced science package	Deployed by long-range rover to form meteorological and seismological network. (Alternatives are penetrators or extended lifetime probes.)

Demonstration represent significant new possible applications. However, a NEP development program must be initiated in the 1980s to be operational in time for a Titan mission around the turn of the century.

The only major alternative propulsion technology is a chemical system using cryogenic liquids (the so-called Orbit Transfer Vehicle or OTV) for Earth escape, followed by gravity assists from Jupiter (in 1998) or from Earth and Venus, followed by aerocapture at Titan in the 2005-2010 time frame.

Main orbiting spacecraft. The principal vehicle for exploration in near-Titan space is an orbiter craft which remains with the NEP system. During the spiral capture process, the spatial structure of fields and particles around Titan can be measured. Following capture, the main spacecraft is parked in a circular polar orbit roughly 600 km above the surface of the body. Such an orbit has relatively little atmospheric drag and is highly desirable for close measurement and deployment of subsidiary system components into the atmosphere and to the surface of Titan.

During operations in near-Titan space, the main spacecraft must support a set of sophisticated remote-sensing instruments needed for global characterization and intensive study. In addition, it must continue to provide essential functions initiated during the interplanetary phases and support for deployed subcraft including navigation and communications with Earth. The estimated data collection volume is estimated at 10^{10} - 10^{11} b/day, significantly greater than the 10^9 b/day characteristic of previous planetary missions. Most of this is accumulated from instruments aboard the main orbiter, with perhaps 10% supplied by subsatellites and surface vehicles. Assuming that all raw data are returned to Earth, the required downlink communications capability is 10^5 - 10^6 b/sec or 3-30 times the Voyager mission capacity from Saturn. However, significant amounts of data compression using advanced machine intelligence techniques should greatly reduce the transmission burden on the terrestrial downlink and also between elements of the Titan Mission.

The technologies developed in present and future planetary missions (especially Galileo, VOIR, and Earth-orbital) are generally applicable to this spacecraft. For instance, while in Titan orbit, the main orbiter is nadir-pointing much like VOIR and many Earth-sensing satellites. Major advancements are expected in the areas of machine intelligence and smart sensors, which suggests an increased capacity for data handling and communications as compared to previous planetary missions by the time of the Titan Demonstration.

Lander/rover. A lander/rover is needed to perform detailed surface and atmospheric measurements as well as the intensive level of study. Deployment of this spacecraft system is deferred until Titan's ground terrain has been fully mapped and an appropriate target site selected.

Atmospheric data are taken during the lander descent phase, and this continues as long as the vehicle remains operational on the planetary surface. Small rocket thrusters are used to guide the craft to a safe place free of large boulders, deep crevasses, or steep slopes. After a soft landing, the surroundings are characterized in preparation for site selection for sample collection. Physical samples are then acquired using extensible manipulators (scoops, drills, slings, etc.) and are immediately analyzed to determine chemical composition, layering effects, evidence for indigenous lifeforms, etc.

After this has been accomplished, the lander requires samples taken from a wider area to complete its preliminary investigations. The general solution to this problem is the rover, a vehicle deployed by the lander and used to explore the local neighborhood and to bring back samples. The simplest rover design might operate no more than 100 m from the lander and would remain almost totally dependent upon it. Such a machine is useful for collecting samples more free of contamination and more representative of the surface than those taken nearer the landing site. However, the Space Exploration Team prefers a more ambitious design, an autonomous rover able to operate up to 10 km from the lander. This larger-area capability permits the lander/rover system to return data which better contributes to an overall understanding of the geological structures of complex sites. Such advanced rovers already have been considered for lunar and martian applications.

It is also necessary to provide the capability of performing intensive studies at more than one surface landing site. This flexibility is possible by deploying multiple lander/rover teams which may be carried from site to site using powered air vehicles for very-long-distance transport. Physical samples could also be returned to stationary landers by similar means. Another possibility is a highly sophisticated long-range rover having a complete set of instruments and sample collection and analysis equipment, and designed for higher speeds, longer traverses (more than 100 km), and enhanced survivability over more difficult terrain with more challenging obstacles. Long-range rovers could visit any number of distinct geologic regions during their lifetimes and might be used to deploy a network of stationary science packages across the surface of the entire planet. The orbit of the main spacecraft is such as to permit regular contact with surface vehicles twice each Titan day (once each Earth week).

The lander/rover system needs extensive machine intelligence capability. Technology requirements are greatest for a long-range rover operating independently in the absence of continuous communications with the main orbiting spacecraft or with Earth. This capability is highly desirable, since without it the operational demands placed on other mission elements — such as the subsatellites for ground-to-orbit Titan uplink or powered air vehicles necessary for

sample and system component transport -- rapidly may become unmanageable.

A significant heritage may be expected from experience gained with the Viking landers and from any future martian or lunar missions, several of which might be approved and flown prior to the Titan Demonstration. One potential major difference is the unknown character of the surface including the possible existence of open liquids on Titan. If fluidic features are widespread it may be necessary to devise new methods of surface mobility and long-distance planetary exploration. New rover concepts for the reduction of machine intelligence requirements by decreased susceptibility to hazards should also be investigated.

Subsatellites. In addition to the main orbiter, subsatellites may be needed for certain specific purposes. One example is a free-flying spacecraft stationed at the L1 Lagrangian point between Titan and Saturn. This could be used to monitor the particle/field environment beyond Titan's magnetosphere, to observe the target atmosphere, and to communicate with mission elements located on the Saturn side of Titan. Another example is a tethered subsatellite system operating within 100 km of the main orbiter -- such multiple devices can more easily distinguish spatial and temporal variations in particles and fields and probe the upper atmosphere (which would cause unacceptable drag on the main spacecraft if it attempted these measurements directly).

The subsatellite concept is new to planetary mission planning. However, these devices currently are projected for use on the Space Shuttle and also are under consideration in connection with manned and unmanned orbital platforms. This technology should become available by the time of the Titan Demonstration (e.g., the spin-stabilization of Mission relay subsatellites). There may also exist some commonality with previous planetary missions such as Pioneer 10/11 and Pioneer Venus.

Atmospheric probes. Several mission components must be sent into Titan's atmosphere at selected locations to make in situ measurements of the air and to carry small instrument packages to the surface. These probes are deployed by the main orbiter from its 600-km circular polar orbit, thus permitting considerable flexibility in choice of geographical entry points and timing. Atmospheric entry probes measure vertical profiles of the atmosphere at the time of deployment, and provide sufficient information to meet mission objectives at the "exploration" level. The Pioneer Venus, Galileo, and proposed Saturn Orbiter Dual Probe (SOP²) missions all include atmospheric entry probes among their equipment.

One large entry probe and at least three small probes are necessary to fulfill the major objectives of Titan exploration. As in the Pioneer Venus mission, all probes measure atmospheric structure, pressure, temperature, etc., whereas

only the large probe takes more detailed data regarding composition, cloud structure, and planetary heat balance. (The large probe considered for the Titan Demonstration is roughly the same size and complexity as the device proposed for the SOP² mission.) Both types of probes also may serve as limited-purpose surface stations.

Powered air vehicles. Many options exist for intensive atmospheric investigation using still more sophisticated vehicles. A superpressure or passive hot-air (Montgolfier) balloon can be designed to float along an isobar for extended periods of time, providing a continuous record of wind speeds and other atmospheric data. Tethered balloons or kites could be used to sample the aerial environment surrounding a surface station. Powered air vehicles such as airplanes, helicopters, and dirigibles can study still larger regions of the atmosphere.

Of the options considered, the powered air vehicle - especially one having an inexhaustible energy supply for long-term operation - appears preferable. Such craft could be used to support extended surface operations, to conduct remote-sensing observations near the base, and even to help collect samples to be returned to the base site for detailed analysis. Regardless of whether the vehicle is an airplane or dirigible, it is highly unlikely that much previous experience will have been acquired with such systems in planetary missions. While the aerodynamic properties of fliers may match those of some Earth-based machines, control and propulsion requirements are likely to differ significantly. Control problems perhaps may be solved using a combination of smart sensors and an advanced machine intelligence capability, together with a satisfactory energy source such as a 10 kW nuclear-power generator to drive an efficient propeller. Titan's atmosphere possibly could be utilized for the production of propellants or buoyant gas.

Packaging the entire system and deploying it at Titan is an additional concern.

Surface science network. A scientific network should be established consisting of at least three permanent sites on the Titanian surface. The network collects seismographic and meteorological data needed to infer subsurface structure and global atmospheric circulation patterns. There are several ways to establish a network, such as (1) using long-range rovers to deploy stationary science packages, (2) deploying surface penetrators dropped from the main orbiter, and (3) extending the lifetime of the atmospheric probes (also dispatched from the main orbiter).

The network concept emphasizes long-term observation - as much as 5 years or more - on Titan's surface. Assuming network stations communicate directly to the main orbiting spacecraft, data must be stored for about a week following collection before uplinking. Each station must be able to function in an extremely cold thermal environment (about 100 K) with internal parts maintained

at reasonable operating temperatures not below 220 K. Stations must be well-coupled to the planetary surface for seismometric purposes but must not thaw crustal ices. One solution is the radiation of excess heat up into the atmosphere.

All of the above components are relatively simple systems, mostly achievable using current or foreseeable aeronautical technology.

3.2.4 Machine Intelligence and Automation Requirements

In outlining the operational mission stages for a Titan demonstration and for the exploration of deep space, a number of automation technology drivers were identified in each of two general categories of system functions:

(1) Mission integrity, including self-maintenance, survival of the craft, and optimal sequencing of scientific study tasks.

(2) Scientific investigation, including data processing and the methodical formation of hypotheses and theories. Both categories impose considerable strain on current AI technology for development in several overlapping areas of machine intelligence. These requirements represent research needs in domains of present concern in the AI community, as well as new research directions which have not yet been taken.

Success in mission integrity (fig. 3.4) requires the application of sophisticated new machine intelligence techniques in computer perception and pattern recognition for imaging and low-level classification of data. This also presupposes the utilization of a variety of remote- and near-sensing equipment. Onboard processing of collected data serves to coordinate the distributed systems and planning activity in terms of reasoning, action synthesis, and manipulation. More capable remote sensing is the key to efficient exploration, making more selective and efficient use of highly complex equipment for atmospheric and planetary surface monitoring.

With respect to reasoning, automated decisionmaking emerges as an important research area. Within this field, development might depart from current expert systems with advancements coming in the form of interacting simulation models of the processes which structure given domains and hypothesis formulating logics. New research directions lie in the areas of alternative computer logics, self-constructing knowledge bases, and self-learning systems.

A need has been identified regarding action synthesis, or procedural sequencing, for representing the relationship between predefined goal states and the current state, and for reducing the discrepancy between the two through automated implementation of subgoals and tasks. Such a system implies the utilization of a sequential informational feedback loop. A more difficult problem is simultaneous

coordination through anticipation, or prediction of the most appropriate action patterns followed by implementation of such action before a large discrepancy occurs. Complementary to the above capability is the capacity for automated construction of unprogrammed goal states as the result of environmental feedback. These latter two technology drivers fall under the general heading of automated learning and are not part of current research interests in the AI community at large.

Another broad technology requirement within the category of mission integrity is manipulation. A fully autonomous system should be capable of self-maintenance and repair, as well as sample collection for data analysis and utilization in decisionmaking processes. The former task presupposes some initial ability for self-diagnosis, while both tasks require a variety of effector capabilities for dealing with a wide range of situational demands. Here, advances in robotics with respect to hand-eye coordination

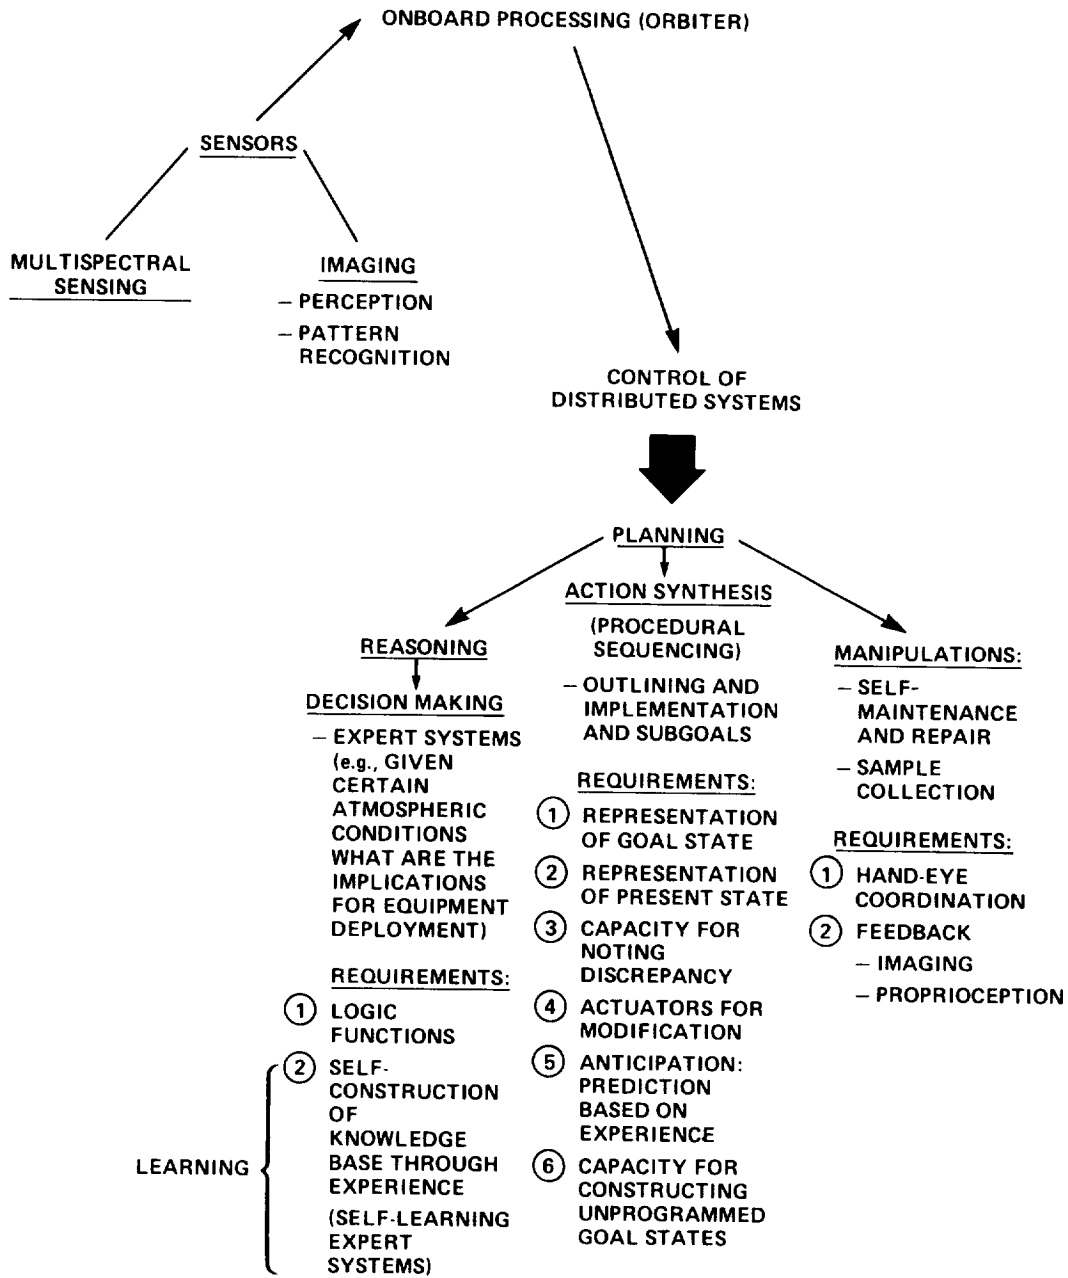


Figure 3.4. – Mission integrity.

and force/proprioceptive feedback systems emerge as significant.

The technology drivers identified for the scientific investigation category of mission functions (fig. 3.5) overlap to some degree those outlined for mission integrity. Automated intelligent planning is perceived as a general requirement in terms of defining scientific goals (both preprogrammed and self-generated) and for the definition of appropriate subgoals. Advanced decisionmaking also is an essential prerequisite for implementing scientific research and for conducting experiments. Decisions such as whether or not an experiment should be carried out, or where and when it should be conducted, probably could be accomplished (as with mission integrity) through extensions of current expert systems technology.

Reduction of collected sensory data to informational categories is yet another significant technology driver. A number of requirements emerge, starting with the ability to describe data at the simplest perceptual level. A higher-order task is the addition of data descriptions to a knowledge base for purposes of classification. This classification may be accomplished in terms of given categories of knowledge requiring some low-level hypothesis generation and testing. More advanced is the necessary capability for reorganizing old categories into new schemes or structures as a consequence of active information acquisition. Under-

lying this form of classificatory activity is again the self-learning process of hypothesis formation and testing. Each of the aforementioned tasks require varying levels of research and development to transform them into fully realized capabilities.

Finally, a requirement exists within the area of communication – transmitting acquired information back to human users. Here the emphasis is on automated selection processes in which an advanced decisionmaking system determines what information and which hypotheses are appropriate and sufficiently interesting to report. The obvious need to communicate with human beings in this case underscores the need for further developments in the field of natural language interfaces.

A scenario illustrating the great complexity of data processing and high-level hypothesis formation capability required for scientific investigation by an autonomous exploration system is presented in appendix 3C.

3.3 Machine Intelligence in Space Exploration Missions

The advanced machine intelligence requirements for general-purpose space exploration systems can be summarized largely in terms of two tasks: (1) Learn new environments, and (2) formulate new hypotheses about them.

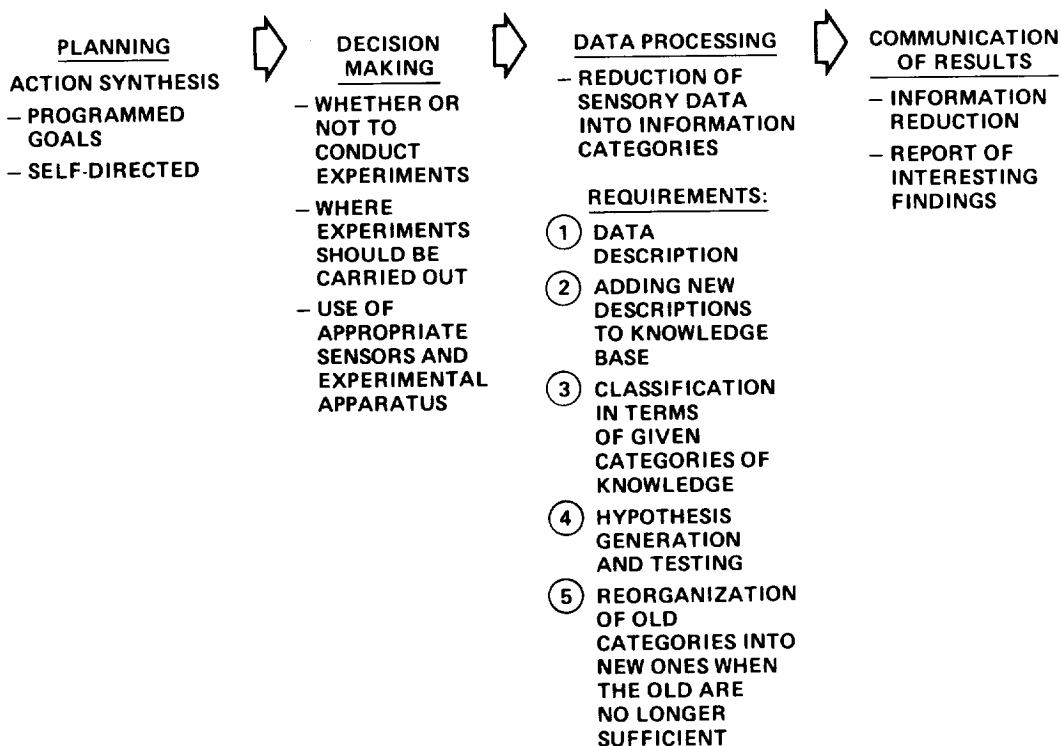


Figure 3.5. – Scientific investigation.

Hypothesis formation and learning have emerged as central problems in machine intelligence, representing perhaps the primary technological prerequisites for automated deep space exploration.

The Titan, outer planet, and interstellar missions discussed by the Space Exploration Team require a machine intelligence system able to autonomously conduct intensive studies of extraterrestrial objects. The artificial intelligence capacity supporting these missions must be adequate to the goal of producing scientific knowledge regarding previously unknown objects. Since the production of scientific knowledge is a high-level intelligence capability, the AI needs of the missions may be defined as “advanced-intelligence machine intelligence,” or, more briefly, “advanced machine intelligence.”

3.3.1 *A Working Definition of Intelligence*

Before an advanced machine intelligence system can be developed and implemented, the concept must be precisely defined and translated into operational terms. One way of doing this is to specify the patterns of inference which constitute the high-level intelligence – the design goal for advanced AI systems. Optimally, designers would have at their disposal an ideal definition of “intelligence” stating the necessary and sufficient conditions for achieving their goal. Such a definition, in addition to precisely stating what intelligence is, also would provide a set of criteria with which to decide the question: “Does entity X possess intelligence?” Unfortunately, no generally accepted ideal definition of intelligence is yet available.

However, a working definition sufficient for the purposes of the present investigation can be formulated. This inquiry addresses the general question of the characteristics of an advanced machine intelligence system needed for autonomous space exploration missions. As such, the investigation should address two questions in particular: “What intelligence capabilities must be designed into space exploration systems?” and “By what criteria will it be determined whether or not the final system actually possesses the high-level intelligence required for the mission?”

American Pragmatism, the major school of American philosophy, developed an account of intelligence that contains the key to a useful working definition (Davis, 1972; Dewey, 1929, 1938; Fann, 1970; Mead, 1934, 1938; Miller, 1973; Peirce, 1960, 1966; and Thayer, 1968). The major figures of this school – John Dewey, William James, George Herbert Mead, and Charles Sanders Peirce – claimed that an entity’s intelligence consists of its ability to reduce the complexity and variety of the world to patterns of order sufficient to support successful action by that entity. For example, human beings have reduced their welter of sensations to patterns of order, e.g., in comparative distinctions between nutrients and non-nutrients, chemical qualitative analysis schemes, and abstract aesthetic concepts. These

patterns are, in turn, the bases of human actions including (following the above examples) satisfaction of the need for food, identification of an unknown chemical compound, and the creation of a work of art.

The Pragmatists further claimed that these action-related patterns of order exhaust an entity’s knowledge. In other words, all knowledge is action-related – indeed, according to Peirce, “to have a belief is to be prepared to act in a certain way.” This view is summarized in the fundamental Pragmatist principle that intelligence is always displayed in action and can be detected only in action. In this view intelligence is a dynamic process, rather than a static state, having at least two dimensions. First, unless an entity has a continuing history of action its intelligence is not displayed, cannot be detected, and therefore cannot be presumed to exist. Second, since a given pattern of order is linked to a related type of action, the success or failure of a particular action reflects on the “correctness” of the underlying pattern of order. An entity can have a continuing history of successful activity only if it can modify or replace those patterns of order which lead to failure. Therefore, an entity’s intelligence is far more than merely the possession of a fixed stock of knowledge – even when this knowledge consists of action-related patterns of order. Rather, intelligence is the ability to preserve a high ratio of successful to unsuccessful outcomes.

The Pragmatists’ account of intelligence can be summarized by this definition: Intelligence is the ability to formulate and revise patterns of order, as evidenced by the eventual emergence of successful over unsuccessful actions. There may well be aspects of intelligence that escape the definition, but nevertheless it provides a useful framework for the present investigation. This is because it focuses on capabilities which must be designed into advanced machine intelligence systems required for autonomous space exploration, as well as on the criteria with which to test for the presence of these capabilities.

A working definition of “advanced machine intelligence” in the context of autonomous scientific investigation of extraterrestrial objects can be formulated by utilizing the above general definition. The Pragmatists held that intelligence is a matter of degree and that among biological entities the question is never intelligence versus nonintelligence, but rather the level thereof. The actions by which biological entities display intelligence range from the amoeba’s avoidance of toxic materials to the human’s acquisition of scientific knowledge. The patterns of order underlying this spectrum of activity are characterized by a wide range of complexity paralleling that of the related actions. Machine intelligence also admits of degrees. Applying the Pragmatists’ general definition is primarily a matter of specifying the level of capabilities with which the investigation is concerned.

In particular, application of the general definition to AI in space applications requires interpreting “actions” to

mean “scientific investigation and mission survival” (the two most complex sets of tasks facing an autonomous exploratory system) and “patterns of order” to mean “the complex abstractive and conceptual structures related to scientific investigation and mission survival” (e.g., hierarchical schemes and terrain maps, respectively). Hence the working definition of advanced machine intelligence in the context of the present study may be summarized as follows:

Advanced machine intelligence is the ability of a machine system to autonomously formulate and to revise the patterns of order required for it to conduct scientific investigations and to survive, as evidenced by continued systemic survival and investigatory behavior despite any environmental challenges it may encounter.

This working definition provides ready answers to the capabilities and criteria issues raised earlier. These responses may be restated from the above definition as follows: (1) An advanced machine intelligence system for autonomous space exploration must possess the capability to utilize already formulated patterns of order and to devise new or revise existing patterns of order; and (2) the criteria by which to determine whether a system actually possesses intelligence is its observed ability to self-correct unsuccessful actions and eventually to act successfully in situations novel to the system.

3.3.2 A Systems Approach to Intelligence

Systems analysis may be used to translate the above definition into practice. Stated in general terms, the design goal is to achieve a machine intelligence capability to autonomously conduct scientific investigations and ensure mission survival. “Intelligence” can be an omnibus term which refers to a broad range of abilities including “knowing,” “emoting,” “fantasizing,” etc. However, only rational cognition such as “knowing” is immediately relevant to machine intelligence for space exploration.

Of course, “knowing” is itself an omnibus term having a range of usages differing somewhat in meaning. In the present context it refers to the rational dimensions of intelligence, the processes of acquiring justified, though possibly fallible, statements about the world and its constituents. Among those dimensions are (1) identifying things and processes, (2) problem-solving, and (3) planning, since the outcomes of each of these processes are statements about the world selected from among a number of alternatives and justified on some basis. The essence of “knowing” in the context of a given environment is the ability to organize and thereby reduce the complexity and variety of perceived events, entities, and processes in the surroundings – a broad general class of rational activity required for machine-intelligent space exploration systems.

A “classification scheme” is any distinction or set of distinctions which can be used to divide events, entities, or

processes into separate classes. By this measure taxonomies, analytical identification procedures, scientific laws and theories (e.g., “ $F = ma$ ” names, hence, distinguishes forces and masses), decision criteria (e.g., go/no-go configurations in a given context), and concepts (e.g., “true” divides all statements into two separate classes) all are examples of classification schemes. Thus, a scheme is any statement, theory, model, formula, taxonomy, concept, categorization, classification, or other representational or linguistic structure which identifies the recurring characteristics of particular environments.

Tasks by which knowing is accomplished may be divided into two distinct types: (1) Utilization of preformulated fixed classification schemes, and (2) generation of new classification schemes or revision of old ones by formulating new components for the schemes. These two task types differ fundamentally both in the characteristics of the tasks and in the types of inference which underlie them.

When preformulated, fixed classification schemes are used, outcomes include identifications, classifications, and descriptions of events, entities, and processes occurring in the environment. These outcomes take the form of statements of the following general types:

- “X is an entity of type A.”
- “Y is an instance of process B.”
- “Z is a class-C event.”

In each case, perceived constituents of the environment are matched with the general classes of constituents into which the classification schemes divide the world. The pattern of inference underlying this type of task is the analytic comparison of actual environmental constituents with “known” assertions about general environmental characteristics. Thus, an important aspect of the utilization of classification schemes is the confrontation of these schemes with the facts of experience. Knowing of this type cannot be successful – indeed, cannot even continue – if the actual state of affairs in the environment and that postulated by the classification schemes differ significantly. So, while the utilization of preformulated classes is an important type of knowing activity, the actual knowing of a given environment is deficient if the schemes are incomplete or incorrect. Knowing can be complete only when new classification schemes can be formulated and incorrect ones revised.

The creation and revision of classification schemes is the second major type of task involved in knowing. The outcomes of this task are either new classification schemes or new parts for preformulated ones. This task can, in turn, be divided into subtasks – the invention of new or revised classification schemes and the testing of these schemes for completeness and correctness prior to general use. Quite different types of inference underlie these two subtasks. Testing new or revised classification schemes requires analytic comparison of the claims made by these schemes with the facts of the world, exactly the same kind of process

involved in the utilization of classification schemes. However, the invention of new or revised schemes demands completely different types of inference. Two patterns of inference comprise this advanced activity – “induction” (included in all standard accounts of inference) and “abduction” (first described by Peirce, 1960, 1966; see also Burks, 1946; Fann, 1970; and Frankfurt, 1958) – as discussed at length below.

The systems approach leads to two important conclusions about machine intelligence (MI). First, MI involves the ability to utilize existing knowledge structures and to invent new ones. Second, although the utilization and invention of classification schemes require the formation of hypotheses, the inference for formulating hypotheses which apply existing classification schemes are logically distinct from the inferences used in formulating hypotheses which invent new or revised classification schemes (see fig. 3.6).

These conclusions have implications for machine intelligence systems designed for autonomous deep space exploration. If classification schemes applicable to the Earth were complete and correct for all extraterrestrial bodies, then an autonomous system utilizing these schemes via analytic inferences alone could successfully complete the knowing process. However, it is probably true that at least some of the available classification schemes are either incomplete or incorrect in the extraterrestrial context and, in any case, the most prudent design philosophy for a space

exploration system would be to assume that gaps do exist. Under the assumption that novelty will be encountered in space, an autonomous exploratory system may successfully complete the knowing process only if it can utilize preformulated classification schemes and also invent new or revised ones, that is, only if it can make inferences of the inductive and abductive types in addition to inferences of the analytic type.

3.3.3 Patterns of Inference for Hypothesis Formation

Analytic, inductive, and abductive inferences will now be characterized in terms of the information inputs and outputs of each. An existence argument for abductive inference, which also establishes its centrality to scientific investigation, is offered, and the process involved in abduction is characterized in some detail. Finally, the requisite state of development for each of the three basic inferential types is contrasted with AI state of the art in the context of autonomous scientific investigation, the ultimate goal.

Analytic inferences are logical patterns by which existing scientific classification schemes (principles, laws, theories, and concepts) are applied to information about the events and processes of the world for the purpose of producing identifications and descriptions of these events and processes as well as predictions and explanations about them (Alexander, 1963; Harré, 1960; Hempel, 1965, 1966;

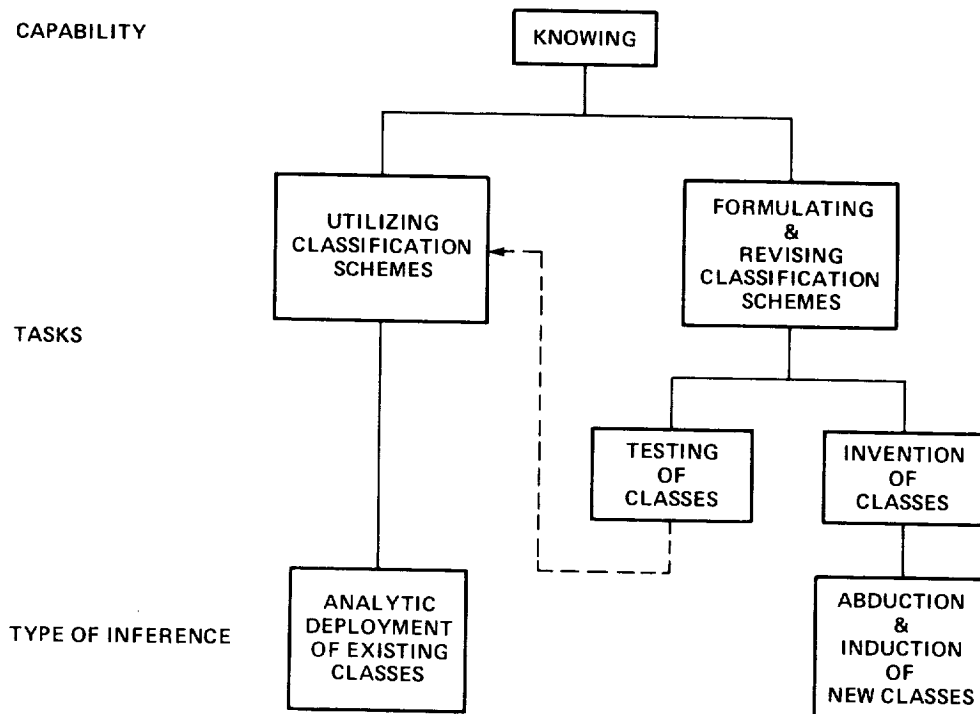


Figure 3.6. – Systems graph for machine intelligence.

Popper, 1963; Wisdom, 1952). This information itself is produced by applying current scientific classification schemes to raw data in an attempt to structure and interpret it. The reasoning is deduction, whether formal deductive logic or other deductivist analytical procedures. Models play an important, though indirect, role in analytic inference (Hanson, 1958; Kuhn, 1970; Toulmin, 1960). The quantitative and symbolic information and the identifications, descriptions, predictions, and explanations which are the outputs of analytic inferences are derived from detailed knowledge such as equations, formulas, laws, and theories. However, standing behind this detailed knowledge is a fundamental model of the "deep structure" of the world which, in effect, provides a rationale for applying that particular kind of detailed knowledge to that specific data. For instance, the kinetic-molecular theory of gases is one such fundamental model whose scientific function is to provide a rationale for searching and then applying a particular kind of detailed knowledge about gases. Figure 3.7 shows the input/output structure of analytic inference.

Inductive inferences are logical patterns for moving from quantitative or symbolic information about a restricted portion of a domain to universal statements about the entire domain (Cohen, 1970; Good, 1977; Hilpinen, 1968; Horton, 1973; Lehrer, 1957, 1970; Rescher, 1961; Salmon, 1967; Skyrms, 1966). There are two somewhat different aspects of inductive inference: Inductive generalization and abstraction. In inductive generalization, some finite set of measurements of an independent variable and its dependent variable are generalized into a mathematical function which

holds for all possible values of those variables. Alternatively, in abstraction, some finite set of symbolic representations of just a few members of some domain is the basis for inferring some abstractive characteristic common to all members of the domain. Examples of abstraction include moving from a set of white objects to the concept of "white," and inferring from the information that all observed ravens are black; the principle that being black is a defining characteristic of ravens. As was the case with analytic inferences, models play an important though indirect role (Hanson, 1958; Kuhn, 1970; Toulmin, 1960). These models serve to restrict the range of mathematical functions or abstractive concepts that can characterize a domain, hence, they focus the inductive inference from information to generalization. For instance, we know that Robert Boyle was guided in the processing of pressure and volume data by a model of gases that required volume to decrease while pressure increased (Toulmin, 1961). Figure 3.8 suggests the input/output structure of inductive inference.

Abductive inferences are logical patterns for moving from an input set that includes:

- some theoretical structure T consisting of models, theories, laws, concepts, classification schemes, or some combination of these,
- some prediction P derived from T by means of an analytic inference, and
- some set of quantitative or symbolic data D which contradict P ($D = \text{not-}P$),

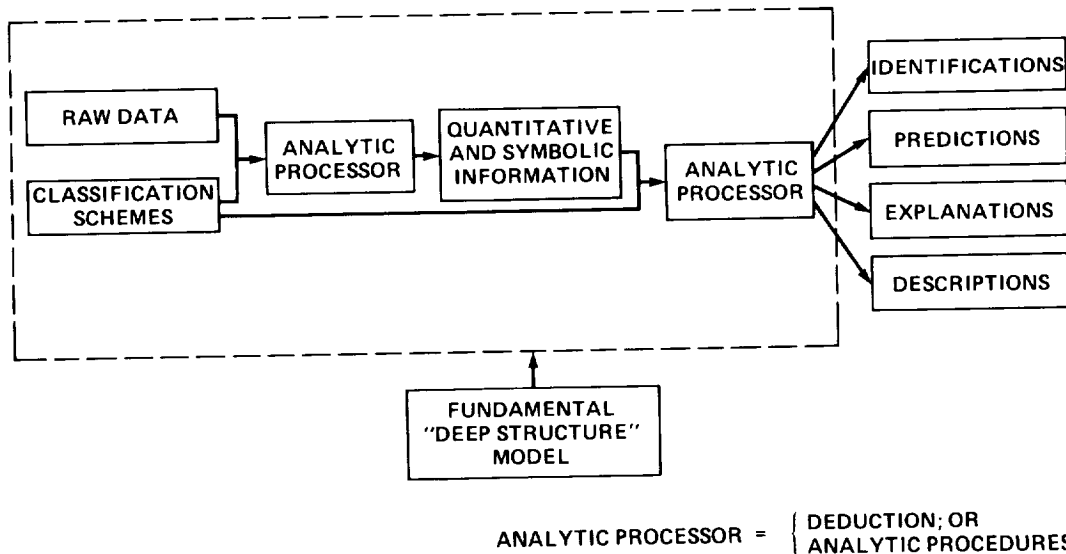


Figure 3.7. – Analytic inference.

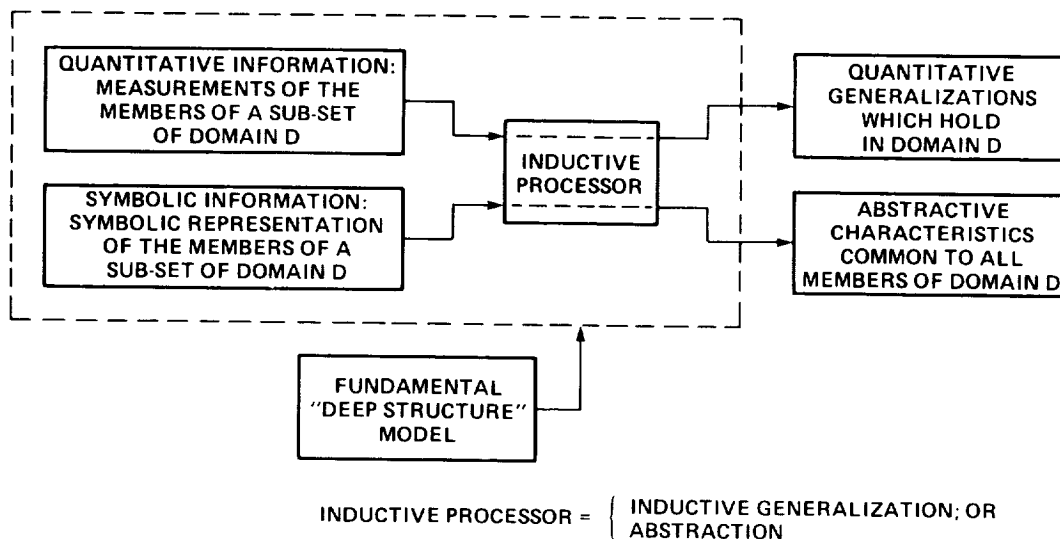


Figure 3.8.— Inductive inference.

to an output set that includes:

- a new or revised theoretical structure T^* ,
- a prediction P^* derived from T^* , and
- a set of quantitative or symbolic data D^* which both agrees with P^* and is the representation of D in T^* ;

that is, $D^* = P^*$ and D^* is the mapping of D onto T^* (Burks, 1946; Davis, 1972; Dewey, 1929, 1938; Fann, 1970; Frankfurt, 1958; Gravander, 1975, 1978; Hanson, 1958, 1961, 1965, 1967, 1969; Kuhn, 1957, 1970, 1977; Lakatos, 1970, 1976, 1977; Mead, 1934, 1938; Miller, 1973; Peirce, 1960, 1966; Simon, 1965; Toulmin, 1960, 1961, 1972; Van Duijn, 1961). Models play a far more important role than in just analytic and inductive inferences: In abduction, fundamental models of processes structuring the world enter directly into the inference. Such models sometimes are the component of a theoretical structure replaced or modified by the inference. Of course, not every replacement or revision of the theoretical structure involves model modification. Those abductive inferences which revise or replace such components as laws or generalizations take the model to be a premise of the inference. The input/output morphology of abductive inference is shown in figure 3.9.

Probably there exists a family of abductive inference species. However, all members of this family must bear many resemblances to one another. Two such family characteristics are particularly important. First, the logical impetus behind the transition from T to T^* is the ability of T^* to explain data which T cannot. Second, the attainment of explanation involves a re-representation of information — i.e., T fails to explain D and T^* explains D^* , where

D^* is not D but rather the representation of D in T^* . As Lakatos (1976) notes, “discovery” is a process in which a theory stated in language L fails to explain a fact; therefore, it cannot adequately be represented in L , so a theory stated in L' must be found to explain it and allow its representation in L' .

Virtually all standard accounts of scientific investigation include analytic and inductive inferences as important components of the logic of science. Abductive inferences are not as widely accepted or understood. Nevertheless, numerous detailed analyses of actual scientific discoveries have demonstrated that there are inferences in these discoveries that are neither analytic nor inductive in nature (Gravander, 1975; Hanson, 1958; Kuhn, 1957; Lakatos, 1977; McMullin, 1978). Examination of these scientific discoveries establishes that the researcher involved in the discovery possessed a determinate set of initial information, including some existing theory and data contradicting a prediction of the theory, and that there is a detailed inference which takes this initial information as its premise and provides the discovery as a conclusion. Whether it is possible to prove that the scientists in question actually followed this inference step by step is irrelevant insofar as the present investigation is concerned. The analyses reported demonstrate the existence of a family of nonanalytic and noninductive inferences which produce new or revised theoretical structures as output. This demonstration constitutes an existence argument for abductive inference.

It cannot be emphasized too strongly that the analysis of actual scientific discoveries is valid only as an existence argument for abduction, not as a research program for mechanizing it. Investigations into the logical process

underlying abductive inference certainly is a first step toward mechanizing the invention of new or revised scientific laws and concepts. But these inferences cannot be demonstrated to be the inference which the scientist followed to the new notion, rather, only an inference having this new notion as its conclusion. Thus, it is not at all clear that a theory of abduction adequate for machine intelligence applications must await a full understanding of human cognition. Quite the contrary; the preferred approach is to attempt to develop a theory of abductive inference on the basis of a direct logical analysis, retreating to the more fundamental problem of human cognition only if the techniques of logical analysis fail.

To consider what might be expected from a direct attack on the logic of abduction a brief characterization of inferential steps constituting such inference is presented below. This characterization takes the viewpoint of some unspecified knower "X," either a scientist or an abductive machine intelligence system.

- (1) X is surprised while using theoretical structure T by some occurrence, O , because O is not among X's set of expectations that are based on T .
- (2) X represents O by a determinate set of data, D .
- (3) X demonstrates that D is more than simply unexpected; it is anomalous in the sense that T predicts not- D .
- (4) X traces not- D back to those components $[T_1, T_2, \dots]$ of its total theoretical structure which entered directly into T 's prediction of not- D .
- (5) X determines which element, T_j , in $[T_1, T_2, \dots]$ is the most likely "villain" behind X's misexpectation.

(6) X attempts to reformulate T_j in such a way that when the new T_j^* is substituted for T_j in a revised T^* , O can be represented by D^* which, in turn, is predicted by T^* . (If successful, the next step is (9) below.)

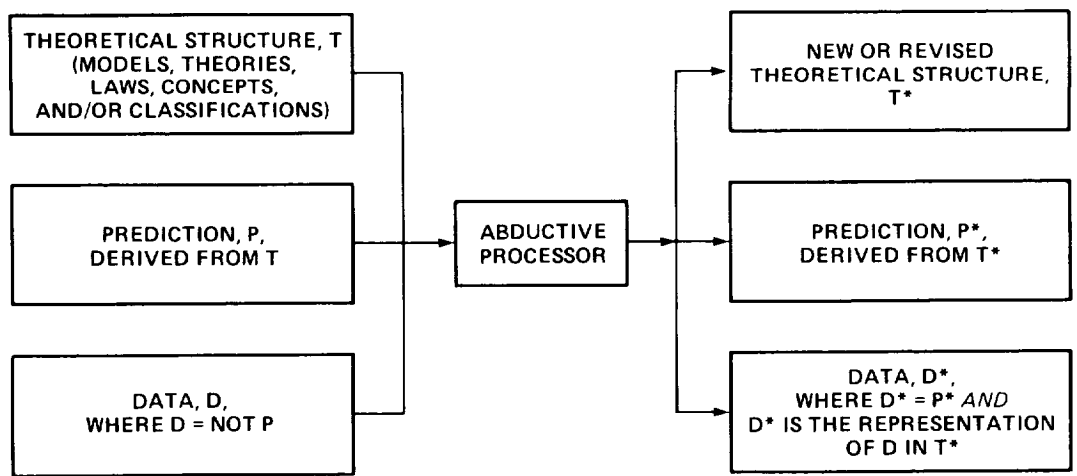
(7) If not successful, X repeats steps (5) and (6) above with the remaining elements of $[T_1, T_2, \dots]$ in order of decreasing likelihood until all possibilities are exhausted. (If successful, the next step is (9) below.)

(8) If still not successful, X repeats steps (5) and (6) with the remaining elements of T , in order of increasing theoretical content and scope, the last component tried being the fundamental "deep structure" model itself.

(9) X makes all adjustments in T^* necessitated by the adoption of T_j^* , including generating a new set of expectations O^* .

(10) X uses T^* until the next "surprising" occurrence.

This characterization of abduction, though not as detailed and precise as that which would result from further investigation, is precise enough to suggest three key problems standing in the way of mechanized abductive inference. First, how should O best be represented as data so that later re-representation is facilitated, and how should these re-representations be performed? Second, is the initial selection of "villains" best achieved by parallel search, hierarchical serial search, or some other technique? Third, can the formulation of the T_j^* replacement of T_j be captured in a stepwise inference in which preceding steps uniquely constrain the selection of the next succeeding step, or must some other technique be used? Note that all of these may be addressed on logical grounds, independent of the broader questions of human cognition.



ABDUCTIVE PROCESSOR = { A FAMILY OF
ABDUCTIVE INFERENCES

Figure 3.9. – Abductive inference.

Finally, it is instructive to contrast state-of-the-art AI treatments of analytic, inductive and abductive inference with the optimal treatment required to achieve working machine intelligence systems with highly advanced capabilities. (See also chapter 6.) First, with respect to analytic inference, current AI research is not addressing the central problem of supporting the detailed knowledge in the classification schemes with fundamental models. Second, although some preliminary work has been done in mechanizing inductive inference (Hajek and Havranek, 1978), this work also has not adequately addressed the basic problem of connecting fundamental models to the generalizing process. Third, only tentative steps have been taken in the development of mechanized abductive inference (Hayes-Roth, 1980), and even these efforts are not grounded on a mature theory of abduction for machine intelligence.

3.3.4 *The Inference Needs of Autonomous Space Exploration Systems*

For an autonomous space exploration system to undertake knowing and learning tasks, it must be capable of mechanically formulating hypotheses using all three of the distinct logical patterns of inference, as follows:

- Analytic inference – needed by the explorer system to process raw data and to identify, describe, predict, and explain events and processes in terms of existing knowledge structures.
- Inductive inference – necessary to formulate quantitative generalizations and to abstract the common features of events and processes, both of which amount to the invention of new knowledge structures.
- Abductive inference – needed by the system to formulate hypotheses about new scientific laws, theories, models, concepts, principles, and classifications. The formulation of this type of hypothesis is the key to the ability to invent a full range of novel knowledge structures required for successful and comprehensive scientific investigation.

Although the three patterns of inferences are distinct and independent, they can be ordered by difficulty and complexity. This ordering is the same as comparing their ability to support the invention of new knowledge structures. Analytic inference is at the low end. An automated system that performs only this type of inference could probably undertake reconnaissance missions successfully. Next is inductive inference. A machine system able to perform this type as well as analytic inference could successfully undertake missions combining reconnaissance and exploration, provided the planet explored is represented well enough by the fundamental models with which the system would be preprogrammed. But if the processes

underlying the phenomena of the new world are not well-represented by the fundamental models, automated combined reconnaissance and exploration missions will require abductive inference. Abduction is at the top of both orderings. It is the most difficult as well as the heart of knowledge invention. An automated system capable of abductive reasoning could successfully undertake missions combining reconnaissance, exploration, and intensive study.

3.3.5 *Cognitive Processes in Intelligent Activity*

One significant technology driver in fully autonomous space exploration is the capacity for learning and the need for adaptive forms of machine intelligence in future space missions (fig. 3.10). However, a review of the literature (Arden, 1980; Boden, 1977; Raphael, 1976) and personal consultations with experts in the field of AI indicate that theoretical and technological research in this area has not seriously been pursued for many years.

For this reason it is useful to approach the goal of adaptive intelligence from the perspective of a related field of study in which it has already received considerable attention: Cognitive psychology. Clearly, descriptions of human thought processes leading to intelligent behavior cannot serve as a direct template for machine intelligence programming – it is a recognized philosophy of the AI community that software need not exactly mimic human processes to achieve an intelligent outcome. Rather, the objective is to describe some aspects of human cognition in hopes of bridging the gap between present limitations in the AI field and the level of machine intelligence likely to be needed in future space exploration missions.

Perception and pattern recognition. The most fundamental kinds of intelligence are perception and the related activity of pattern recognition. Each has been the subject of much study by cognitive and physiological psychologists. For example, evidence from Sperling (1960) suggests that perceptual input is held briefly in a sensory buffer register, thus, permitting the activation of control processes to encode the data in terms of meaningful categories. Stimuli presented to the human sensorium arrive in conscious awareness first as some perceptual-level description, then later with some useful label attached. Exactly how these processes work remains unknown, in part because perception occurs below the subject's level of awareness. Progress to date provides only partially integrated theories of perceptual data handling, yet these are sufficiently well-developed to deserve a brief review in the context of the present study.

A definition of perception at the descriptive level, popular in the psychological literature, holds that sensory processing is essentially inferential or interpretive, based on

raw sensory cues available in the environment, and produces and subsequently tests interpretations about what the world looks like. The percept is the phenomenological result of the interpretation. In this view, perception is a subconscious, "hard-wired" constructive process involving the formation of a hypothesis, a test of that hypothesis, and a consequent decision as to whether the hypothesis accurately encompasses the sensory information. The literature of psychology contains much evidence to support such a description as a reasonable characterization of human perception (Neisser, 1967; Rock, 1975), and the AI community has accepted, in principle, a similar view (Arden, 1980). However, the techniques and operations typically employed to achieve computer pattern-sensing generally fail to properly incorporate the notion of perception and recognition as active constructive processes.

Cognitive psychological theory has largely emphasized two general approaches in characterizing pattern recognition schemes – template matching and feature extraction theory. Each has a different focus of attention with respect

to the three major aspects of recognition called "description," "representation," and "matching" (of new images against stored representations).

Template matching theorists propose that a literal copy of perceived stimuli stored in memory is matched against new incoming stimuli. Although this view has been criticized as too simplistic and naive (Klatsky, 1975; Neisser, 1967), updated versions of the hypothesis still hold sway. For instance, one modification retains the notion that literal copies are stored in memory but suggests that new percepts are "normalized" before matching. In this view, some precomparison processing takes place in which edges are smoothed out, oriented in the appropriate plane, and centered with respect to the surrounding field. In addition, image context helps in the normalizing process by reducing the number of possible patterns the stimulus might match (Klatsky, 1975). In the field of AI technology, the Massively Parallel Processor or "MPP" (an imaging system currently under development at Goddard Space Flight Center) uses visual data-handling techniques with characteristics

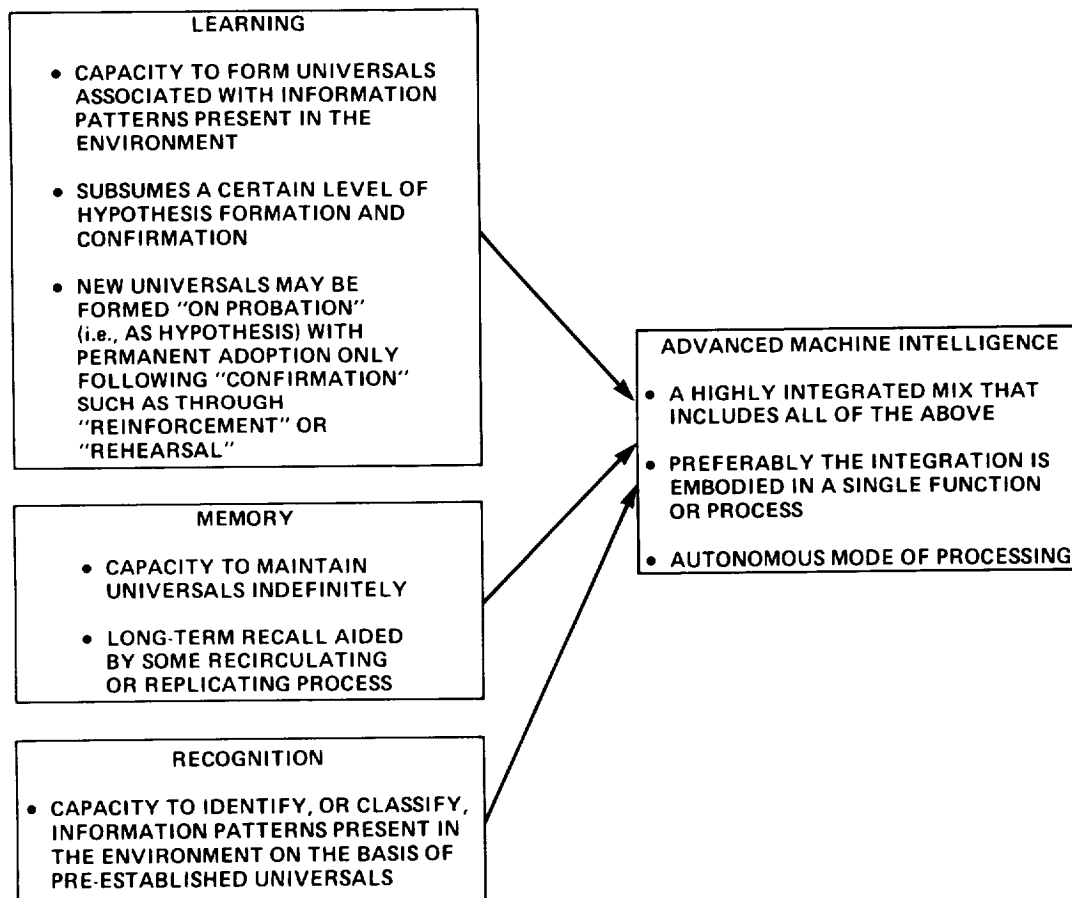


Figure 3.10.— Adaptive machine intelligence for advanced space exploration.

remarkably similar to those described in the normalizing and template matching theories. Given information on its sensory perspective and images stored in its memory, the MPP performs precomparison processing to orient incoming images for compatibility with stored images.

Another hypothesis of perception with similar assumptions is feature detection or feature extraction theory. According to this formulation a pattern may be characterized as a configuration of elements or features which can be broken down into constituent subcomponents and put back together again. Recognition is a comparison process between lists of stored features (which when combined, constitute a pattern) and features extracted from incoming stimuli (Klatsky, 1975). An early theoretical AI model of the feature detection hypothesis was Pandemonium (Selfridge, 1966). This system performs a hierarchical comparison of low-level through higher-order features until the incoming pattern is recognized. More recent scene analysis paradigms have grown from similar assumptions that the raw scene must be "segmented" into regions, or edges of regions, out of which desired objects may be constructed (Arden, 1980; Barrow, private communication, 1980). Scene-analysis models developed on the basis of higher-order features of greater complexity than those proposed by Selfridge have achieved moderate success in limited environments. The major problem is that the system can only deal with familiar or expected input data. All categories within which items are recognized, must be explicitly defined by the programmer in terms of their subcomponents. This eliminates the possibility of recognition processes in novel environments.

Reviewed together, template matching and feature detection reflect the processes modeled by most AI imaging and pattern recognition research. Hence, current AI systems are incapable of handling new category construction and other advanced perceptual tasks which might be required in future space missions. This limitation suggests that an alternative approach to the problem of automated pattern recognition may be needed.

Despite abundant research supporting the existence of feature detectors in humans (Hubel and Wiesel, 1966; Lettvin et al., 1959), other evidence suggests that feature and template theory do not provide a complete explanation of recognition. The above approaches are regarded today as unsophisticated in their conception of how events are mentally represented, and erroneous in ignoring the problem of how representations are achieved. Experiments conducted by Franks and Bransford (1971) indicate that the human mental representation used for feature comparison may be prototypical and holistic rather than literal and elemental. That is, what is actually stored in memory is the product of an active construction, developed over time. In this view the cognitive system extracts and stores the converging "essences" of items to which it is exposed, and this abstraction is then utilized in the recognition process. The empha-

sis is on conceptual representational construction and conceptually driven (top-down) processing, rather than matching and data-driven (bottom-up) processing. The advantage of a prototype approach to perception is that minor distortions or transformations within a limited range will not interfere with the recognition process.

The prototype approach may be considered in terms of two different aspects – the abstract analogical nature of representation and category or concept construction. With respect to machine intelligence, perhaps the closest approximation to the notion of prototypical representation is illustrated by Minsky's "frame" concept. A frame in Minsky's formulation is a data structure for representing a stereotyped situation (Minsky, 1975) and corresponds in many ways to the psychological notion of schema (Bartlett, 1961). Though not really analogical in nature, the frame conception contributes to scene analysis by permitting the system to access data in a top-down fashion and to utilize generalized information without relying on simplistic features. The frames, however, must be described within the system by a programmer and are relatively static. There is no capability for frame reorganization as a result of experience.

Consider now the second aspect of the prototype approach, the construction of abstract categorical representations. Category construction may be viewed as a brand of concept formation. Experimental evidence suggests that the formation of new conceptual categories is the result of a hypothesis generation and testing process (Levine, 1975) in which recursive operations are evoked which infer hypotheses about how a number of particulars are related and then test those hypotheses against feedback information from the environment. Some additional evidence suggests that a number of these hypotheses may be tested simultaneously (Bruner et al., 1956). The result is considered an abstract analogical representation capturing an essence which subsumes all the particulars. Since the hypothesis theory of concept formation typically has been considered in the context of conscious processes, it may seem somewhat far afield of perceptual processing. However, since perception itself has been described as an unconscious inferential process, it may be the case that similar underlying logical operations are at work in the formation of higher-order concepts, prototypes, and in perceptual construction. The precise nature of acquisition, how an "elegant" hypothesis is formed, is not clearly specified in any of these theories. (See section 3.3.3.)

Only a minimal amount of work has been done on AI approaches to the formation of new conceptual structures. A classic attempt was Winston's concept formation program in which a machine was taught through example to acquire new concepts (e.g., the architectural concept of "arch"). Using informational feedback from the programmer as to whether a particular example illustrated the concept or not, and by assessing the essential similarities and differences

among the examples it was shown, Winston's software created structural descriptions of the essentials of the concept in the form of a semantic network.

The function of such a program may appropriately be defined as concept learning. However, the programming techniques appear more closely wedded to the notion of concepts as feature lists rather than as prototypical, analogical structures. This "feature view" has theoretical limits in the domains of human and artificial intelligence since a number of abstract categories can be identified in which constituent members have a few or no structural features in common but whose relationship is either more functional in nature or salient "in more broadly specifiable terms" (Boden, 1977; Rosch and Mervis, 1975). Salience for Winston's program relates only "to categorizations made by its human teacher for human purposes" (Boden, 1977). It is difficult to see how a program with a feature list assumption could move beyond predefined categories to handle the construction of new abstract concepts. This is a significant constraint on state-of-the-art AI technology in terms of future space missions requiring autonomous exploration in novel environments where "there is no guarantee that categorizations previously found useful would still be salient" (Boden, 1977).

Genetic epistemology. One final consideration with respect to intelligent activity comes from Jean Piaget's work on genetic epistemology. This topic is relevant to the issues addressed in this chapter because genetic epistemology offers one of the most comprehensive views of intelligence to be found in the literature today. Piaget's conceptions of the underlying processes of "natural" intelligence encompass the behavioral and cognitive activities of humans and animals. Moreover, the processes are sufficiently general possibly to be captured in a nonliving artifact which would then serve as an effective realization of non-natural intelligence (Piaget, 1970).

How can intelligence be characterized in terms of structures and processes so that it might be embodied in a computer system? One important assumption of Piaget's theory is that any account of the evolution of cognitive activity and intelligence must include the nonteleological aspects of adaptation and purpose. The process of equilibration, a regulative function which propels the subject toward more inclusive and stable interactions with its environment, is basic to the theory. The deterministic result of equilibrium is seen as a characteristic structuring of the relations between subject and environment (Piaget, 1963).

There are two processes that subjects must coordinate in order to achieve a state of equilibrium: Assimilation and accommodation. Assimilation, exhibited by all organisms, is the functional aspect of structure formation by which subjects, acting on their environment, modify it in terms of existing structures (Piaget, 1970). Each organism possesses a set of generalized behavior patterns, or action schemes,

which support its repetitive modification of its environment for the purpose of producing an expanded set of interactions. Accommodation is the modification of the assimilatory cycle itself as a result of the subject's interactions with its surroundings (Piaget, 1963). Accommodation involves the transformation of existing structures in response to continuous environmental stimulation. The result is the construction of new categories of experience which then become part of the organism's general behavioral repertoire.

For Piaget, these "schemes" are the basic units for structuring knowledge (Rosenberg, 1980), the means by which all overt behavioral and cognitive activity is organized. The notion of "scheme" defined by Piaget has certain similarities to Minsky's "frames" as the basic units of knowledge representation. Both notions imply a top-down processing schedule for intelligent activity. However, the two notions differ dramatically in terms of their dynamics. The frame permits a kind of assimilatory activity (organization of particulars within its structure) but the structure itself is relatively static — there seems to be no possibility for reorganization of the structure (the frames) in response to experience. Alternatively, the scheme emphasizes both assimilative and accommodative processes. Accommodation in this case is the restructuring of available schemes into new higher-order schemes which subsume all previous particulars while simultaneously permitting the inclusion of new ones. Again the primary gap between the level of intelligence available with current AI approaches and that which characterizes more advanced intelligent activity appears in the domain of emergent change. Transforming present knowledge structures into new higher-order schemes is a prerequisite for fully intelligent activity, and this capability is absent from state-of-the-art AI techniques.

While the utilization of a genetic epistemological framework has not yet received much study by researchers in the AI field, it has attracted some recent attention in other quarters. For instance, Rosenberg (1980) suggests a number of ways to blend Piaget's theory and current AI methodology to their mutual benefit. Perhaps this represents the beginning of a recognition of the need for comprehensive formulations of natural intelligence to be incorporated into the development of a theory of intelligence in nonhuman artifacts.

3.4 Technology Drivers for Automated Space Exploration

The most important single technology driver for automated space-exploration missions of the future is advanced machine intelligence, especially a sophisticated MI system able to learn new environments and to generate scientific hypotheses using analytic, inductive, and abductive reasoning. Within the AI field the most powerful technology driver is the demonstrable need for an abductive inferential capability useful for inferring new successful knowledge

structures from failed ones. Required machine intelligence technologies include:

- Autonomous processing (essentially no programming)
- Autonomous “dynamic” memory
- Autonomous error-correction
- Inherently parallel processing
- Abductive/dialectic logical capabilities

- General capacity for acquisition and recognition of patterns
- Universal “Turing Machine” computability.

Numerous other supporting technologies also are essential for the staging of autonomous space exploration missions, including low-thrust propulsion systems; general-purpose surface exploration vehicles able to function on both solid and fluid surfaces; reconfigurable sensor nets and smart sensors; flexible, adaptive general-purpose robot manipulators; and distributed intelligence/database systems.

3.5 References

- Alexander, Peter: *A Preface to the Logic of Science*. Sheed & Ward, London, 1963.
- Arden, Bruce W., ed.: *What Can Be Automated? The Computer Science and Engineering Research Study (COSERS)*, MIT Press, 1980. 934 pp.
- Bartlett, Frederic Charles: *Remembering: A Study in Experimental and Social Psychology*. The Cambridge Univ. Press, 1932. Revised Edition, 1961.
- Black, D. C., ed., *Project Orion: A Design Study of a System for Detecting Extrasolar Planets*, NASA SP-436, 1980.
- Boden, Margaret: *Artificial Intelligence and Natural Man*. Basic Books, Inc., 1977.
- Bruner, G. S.; Goodnow, J. J.; and Austin, G. A.: *A Study of Thinking*. Wiley, New York, 1956.
- Burks, Arthur: *Peirce's Theory of Abduction*. *Philosophy of Science*, vol. 13, 1946, pp. 301-306.
- Cassenti, B. N.: *A Comparison of Interstellar Propulsion Systems*, AIAA paper 80-1229, AIAA/SAE/ASME 16th Joint Propulsion Conf., 30 June-2 July, 1980. Hartford, Connecticut.
- Cohen, Laurence Jonathan: *The Implications of Induction*. Methuen Press, London, 1970.
- Davis, William H.: *Peirce's Epistemology*. Martinus Nijhoff, The Hague, 1972.
- Dewey, John: *Experience and Nature*. W. W. Norton and Co., New York, 1929. Revised Ed.
- Dewey, John: *Logic: The Theory of Inquiry*. Henry Holt and Co., New York, 1938.
- Fann, K. T.: *Peirce's Theory of Abduction*. Martinus Nijhoff, The Hague, 1970.
- Frankfurt, Harry G.: *Peirce's Notion of Abduction*. *J. of Philosophy*, vol. 55, July 1958, pp. 593-597.
- Franks, Jeffrey J.; and Bransford, John D.: *Abstraction of Visual Patterns*. *J. of Exper. Psych.*, vol. 90, no. 1, 1971, pp. 65-74.
- Freitas, Robert A., Jr.: *Interstellar Probes: A New Approach to SETI*. *J. British Interplanetary Soc.*, vol. 33, March 1980a, pp. 95-100.
- Freitas, Robert A., Jr.: *A Self-Reproducing Interstellar Probe*. *J. British Interplanetary Soc.*, vol. 33, July 1980b, pp. 251-264.
- Good, I. J.: *Rationality, Evidence, and Induction in Scientific Inference*. E. Elcock and D. Michil, eds., *Machine Intelligence*, Halstead Press, New York, 1977, pp. 171-174.
- Gravander, Jerry Wallace: *Newton's New Theory About Light and Color and the Hypothetico-Deductive Account of Scientific Method: Scientific Practice contra Philosophic Doctrine*. Ph.D. dissertation, Univ. of Texas, 1975. Univ. Microfilms, Ann Arbor, Michigan, 1975.
- Gravander, Jerry Wallace: *Mead's Logic of Discovery*. Michael P. Jones et al., eds., *The Individual and Society*, The Southwestern J. Philosophy Press, Norman, Oklahoma, 1978, pp. 187-207.
- Hajek, Peter; and Havranek, Tomas: *Mechanizing Hypothesis Formation: Mathematical Foundations for a General Theory*. Springer-Verlag, Berlin, 1978.
- Hanson, Norwood Russell: *Patterns of Discovery: An Inquiry into the Conceptual Foundations of Science*. Cambridge University Press, Cambridge, 1958.
- Hanson, Norwood Russell: *Is There a Logic of Discovery?* In H. Feigl and G. Maxwell, eds., *Current Issues in the Philosophy of Science*, Holt Rinehart and Winston, New York, 1961, pp. 20-35.
- Hansen, Norwood Russell: *The Idea of a Logic of Discovery*. *Dialogue*, vol. 4, 1965, pp. 48-61.
- Hanson, Norwood Russell: *An Anatomy of Discovery*, *The Journal of Philosophy*, vol. 64, 1967, pp. 321-352.
- Hanson, Norwood Russell: *Perception and Discovery: An Introduction to Scientific Inquiry*. Freeman, San Francisco, 1969.

- Harré, Romano: *An Introduction to the Logic of the Sciences*. Macmillan, London, 1960.
- Hayes-Roth, Frederick: *Theory-Driven Learning: Proofs and Refutations as a Basis for Concept Discovery*, paper delivered at the Workshop on Machine Learning, Carnegie-Mellon University, July, 1980.
- Hempel, Carl Gustav: *Aspects of Scientific Explanation*. Free Press, New York, 1965.
- Hempel, Carl Gustav: *Philosophy of Natural Science*. Prentice-Hall, Englewood Cliffs, New Jersey, 1966.
- Hilpinen, Risto: *Rules of Acceptance and Inductive Logic*. North-Holland, Amsterdam, 1968.
- Horton, Mary: In *Defense of Francis Bacon: A Criticism of the Critics of the Inductive Method*. *Studies in the History and Philosophy of Science*, vol. 4, 1973, pp. 241-278.
- Hubel, D. H.; and Wiesel, T. N.: *Receptive Fields, Binocular Interaction, Functional Architecture in Cats' Cortex*. L. Uhr, ed., *Pattern Recognition*. John Wiley and Sons, Inc., New York, 1966, pp. 262-277.
- Klatsky, Roberta L.: *Human Memory: Structures and Processes*. W. H. Freeman and Co., San Francisco, 1975.
- Kuhn, Thomas Samuel: *The Copernican Revolution: Planetary Astronomy in the Development of Western Thought*. Harvard Univ. Press, Cambridge, 1957.
- Kuhn, Thomas Samuel: *The Structure of Scientific Revolutions*. (Second ed.), University of Chicago Press, Chicago, 1970.
- Kuhn, Thomas Samuel: *The Essential Tension: Selected Studies in Scientific Tradition and Change*. Univ. of Chicago Press, 1977.
- Lakatos, Imre: *Falsification and the Methodology of Scientific Research Programmes*. In Imre Lakatos and Alan Musgrave, eds., *Criticism and the Growth of Knowledge*. University Press, Cambridge, 1970, pp. 91-196.
- Lakatos, Imre: *The Changing Logic of Scientific Discovery*. Cambridge Univ. Press, Cambridge, 1977.
- Lakatos, Imre: *Proofs and Refutations: The Logic of Mathematical Discovery*. Cambridge Univ. Press, Cambridge, 1976.
- Lehrer, Keith: *Induction, Rational Acceptance, and Minimally Inconsistent Sets*. G. Maxwell and R. L. Anderson, Jr., eds., *Induction, Probability, and Confirmation*, University of Minnesota Press, Minneapolis, 1957, pp. 295-323.
- Lehrer, Keith: *Induction, Reason and Consistency*. *British Journal for the Philosophy of Science*, vol. 21, 1970, pp. 103-114.
- Lettvin, J. W. *et al.*: *What the Frog's Eye Tells the Frog's Brain*. *Proceedings of the IRE*, vol. 47, Nov. 1959, pp. 1940-1951.
- Levine, Marvin J.: *A Cognitive Theory of Learning*. Lawrence Erlbaum Associates, Hillsdale, New Jersey; Halsted Press, 1975.
- Martin, A. R., ed.: *Project Daedalus: The Final Report on the BIS Starship Study*. *British Interplanetary Soc. Suppl.*, British Interplanetary Soc., London, 1978.
- McMullin, Ernan: *The Conception of Science in Galileo's Work*. R. Butts and J. Pitt, eds., *New Perspectives on Galileo*, Reidel, Dordrecht, 1978, pp. 209-257.
- Mead, George Herbert: *The Philosophy of the Act*. Charles W. Morris, John M. Brewster, Albert M. Dunham, and David L. Miller, eds., Univ. of Chicago Press, Chicago, 1938.
- Mead, George Herbert: *Mind, Self, and Society*. Charles W. Morris, ed., Univ. of Chicago Press, Chicago, 1934.
- Miller, David L.: *George Herbert Mead: Self, Language, and the World*. Univ. of Texas Press, Austin, 1973.
- Minsky, M.: *A Framework for Representing Knowledge*. Patrick Henry Winston, ed., *The Psychology of Computer Vision*, McGraw Hill, 1975.
- Neisser, Ulric: *Cognitive Psychology*. Appleton-Century-Crofts, New York, 1967.
- Peirce, Charles Santiago Sanders: *Collected Papers, Volumes 1-6*. Charles Hartshorne and Paul Weiss, eds., The Belknap Press of Harvard University Press, Cambridge, 1960.
- Peirce, Charles Santiago Sanders: *Collected Papers, Volumes 7-8*. Arthur W. Burks, ed., The Belknap Press of Harvard University Press, Cambridge, 1966.
- Piaget, Jean: *The Psychology of Intelligence*. Littlefield, Adams and Co., New Jersey, 1960.
- Piaget, Jean: *Structuralism*. Basic Books, Inc., New York, 1970.
- Popper, Karl Raimund: *Conjectures and Refutations*. Routledge and Kegan Paul, London, 1963.
- Raphael, B.: *The Thinking Computer: Mind Inside Matter*. W. H. Freeman and Co., San Francisco, 1976.
- Rescher, Nicholas: *Rules of Inference and Problems in the Analyses of Inductive Reasoning*. *Synthese*, vol. 13, 1961, pp. 242-251.
- Rock, I.: *An Introduction to Perception*. Macmillan, New York, 1975.
- Rosch, E.; and Mervis, C.: *Family Resemblances: Studies in the Internal Structure of Categories*. *Cognitive Psychology*, vol. 7, 1975, pp. 573-605.
- Rosenberg, J.: *Piaget and Artificial Intelligence*. Published in *Proceedings of the First Annual Conference on Artificial Intelligence*, 1980.
- Salmon, Wesley C.: *The Foundations of Scientific Inference*. Univ. of Pittsburg Press, Pittsburg, 1967.
- Schappell, Roger T., et al.: *Application of Advanced Technology to Space Automation*, NASA CR-158350, 1979.
- Selfridge, O. G.: *Pandemonium: A Paradigm for Learning*. L. Uhr, ed., *Pattern Recognition*. John Wiley and Sons, New York, 1966, pp. 339-348.
- Simon, Herbert A.: *The Logic of Rational Decision*. *The British J. Philosophy Sci.*, vol. 16, 1965, pp. 169-186.

- Skyrms, Brian: *Choice and Chance: An Introduction to Inductive Logic*. Dickenson, Belmont, California, 1966.
- Sperling, George: *The Information Available in Brief Visual Presentations*. *Psychological Monographs*, vol. 74, no. 11, 1960, pp. 1-29.
- Thayer, Horace Standish: *Meaning and Action: A Critical History of Pragmatism*. The Bobbs-Merrill Company, Inc., Indianapolis, 1968.
- Toulmin, Stephen Edelston: *The Philosophy of Science*. Harper and Row, New York, 1960.
- Toulmin, Stephen Edelston: *Foresight and Understanding: An Enquiry into the Aims of Science*. Indiana Univ. Press, Bloomington, 1961.
- Toulmin, Stephen Edelston: *Human Understanding, Volume I*. Clarendon Press, Oxford, 1972.
- Valdes, Francisco; and Freitas, Robert A., Jr.: *Comparison of Reproducing and Non-reproducing Starprobe Strategies for Galactic Exploration*. *J. British Interplanetary Soc.*, vol. 33, Nov. 1980, pp. 402-408.
- Van Duijn, P.: *A Model for Theory Finding in Science*. *Synthese*, vol. 13, 1961, pp. 61-67.
- Wisdom, John Oulton: *Foundations of Inference in Natural Science*. Methuen, London, 1952.

APPENDIX 3A

WHY INTERSTELLAR SPACE EXPLORATION?

The first question a skeptic today might ask is: "Why an interstellar mission?" (fig. 3-11). Twenty years ago many people similarly inquired "Why go to the Moon?" Besides political reasons, there were other goals when the Apollo Program began. For instance, scientists had high hopes for a better understanding of the Earth, the Moon, and the Universe. Yet, although the Solar System is many worlds with countless strange phenomena, still its scientific treasures are miniscule in comparison to those of the Galaxy. Interplanetary space travel is no longer a dream, but a reality — the new dream is interstellar space travel.

Mankind cannot survive forever tied to the cradle of the Earth. In perhaps six billion years our Sun will burn itself out, exhausted of its thermonuclear fuel. But Earth should become uninhabitable long before that. Nuclear

war, asteroid collisions, or innumerable other planet-scale disasters could wipe out much of terrestrial life including mankind. The human species remains at risk until humanity extends itself beyond its homeworld. As a young person eventually must leave his parents' home to seek his own path, so must mankind extend its grasp far beyond its ancestral birthplace. Interstellar travel offers the hope of ultimate long-term perpetuation of human life.

We, as a species, possess a deep instinct to survive. Adventure and risk attract many people. It is possible to imagine a manned interstellar mission with all of the above in mind and more, and to dream of life afresh on an alien world with room to grow and a chance for countless new beginnings. Is this really so different from the early settlers who crossed the Atlantic in search of a "New World?"




Figure 3.11.— A spacecraft sent out to the stars to discover and explore new worlds.

APPENDIX 3B

EXCERPT FROM HYPOTHETICAL TITAN MISSION LOG

1. Central Spacecraft Computer.
 - a. Monitors progress of all operations.
 - b. Initializes all tasks.
 - c. Keeps log and communicates to Earth.
 - d. Makes and tests hypotheses when anomalies from predefined "Planet Model" are found.
2. Spacecraft Imager.
 - a. Records images in "snapshot" fashion (on retina-like array of detectors).
 - b. Finds features asked for by central spacecraft computer.
 - c. Notes anomalies in images from predefined "Planet Model."
 - d. Is capable of describing images in terms of predetermined concepts.
 - e. Updates "Planet Model" with new information based on image input.
3. Spacecraft Control Image Processor.
 - a. Processes data scanned by visible, IR, microwave, and other image sensors. Puts data into "Planet Model."
 - b. Performs tests asked for by Central Spacecraft Computer on this data.
 - c. Identifies surface features and matches features to concepts stored in "Planet Model."
 - d. Updates "Planet Model" based on new information.
4. Lander Central Image Processor.
 - a. Main lander vision processor capable of looking in any direction.
 - b. Performs scene analysis to locate objects of interest on surface and to locate position of lander.
 - c. Retinal-type sensor input.
 - d. Adds surface data to "Planet Model."
5. Lander Guidance Image Processor.
 - a. Processes image data to determine safe path from present location to assigned destination.
 - b. Updates "Planetary Model" contour map.
 - c. Finds obstacles on ground during descent.
6. Central Lander Computer.
 - a. Handles requests from central spacecraft computer.
 - b. Plans lander actions based on these requests.

- c. Assigns tasks to Lander Central Image Processor and Lander Guidance Image Processor.

HYPOTHETICAL SHIP'S LOG – SPACESHIP TITAN July 4, 2010

REPORT: CENTRAL SPACECRAFT COMPUTER (CSCC)

- 9:00 am GMT – Have Titan in view on spacecraft imager. Based on size of disk at 300 mm focal length, we are 504300 km from satellite. This agrees ± 100 km with microwave (5680 GHz) ranging system.
- 9:10 am GMT – Zoomed to 3000 mm focal length. Approximately 5:30 am local zero meridian time.

OBSERVATIONS – SPACECRAFT IMAGER.

1. Satellite generally dark in appearance with some light blotches. Surface appears smoother than Earth's moon.
2. No polar caps observed.
3. Terminator sharp.
4. Limb at equator bright at surface.

HYPOTHESIS: CSCC – This indicates high-density model of atmosphere to be correct.

QUERY: CSCC – High-density atmosphere does not agree with sharp terminator observed.

CSCC TO SPACECRAFT IMAGER:

- Task 1: When at range 20000 km, observe limb using spectral analysis procedure.
- Task 2: Measure spectral reflectance over 200 km \times 500 km area centered on terminator at equator.

OBSERVATIONS – SPACECRAFT IMAGER (cont.).

5. Several dark areas on surface near limb at 40° north latitude. Perpendicular projection would show these to be roughly circular. Areas very dark in center but lightens (like an inverse conic function) toward edges.

CSCC TO SPACECRAFT IMAGE PROCESSOR:

Task 3: When in orbit, correlate contour map, spectral data, and geological analysis over this area to test for possible volcanic activity.

July 10, 2010

REPORT: CSCC

1:30 pm GMT – Disk size at 300 mm focal length indicates range 20100 km from satellite. Agrees ± 100 km with microwave ranging system.

1:35 pm GMT – Terminator observed through 560 nm, 630 nm, and 1080 nm spectral bands. Relative average reflectance 0.2, 1.5, 0.1, respectively, indicating dense atmosphere refraction properties.

QUERY: CSCC – Since dense atmosphere hypothesis was reinforced, why sharp terminator?

CSCC TO SPACECRAFT IMAGER:

Task 4: Cancel Task 1.

Task 5: Measure geological activity in this area during orbit.

July 20, 2010

REPORT: CSCC

7:55 am GMT – Orbit insert completed at 6:00 am. Beginning Titan resource survey for Planet Model. Visible, IR, and microwave scanners operating. Data being processed and stored. Contour map being generated. Gravity anomaly experiment initialized.

8:00 am GMT – OBSERVATION – SPACECRAFT IMAGER

1. White haze in image slightly obscuring surface. Cloud-like rather than bright surface.

CSCC TO SPACECRAFT IMAGE PROCESSOR:

Task 6: Check Planet Model stereo imagery to verify haze above surface.

Response: Spacecraft Central Image Processor.
1. Verifies, white haze is above surface.

Task 7: Check stored spectral data to estimate spectral characteristics of haze.

Response: Spacecraft Central Image Processor.
1. Spectrum of samples areas matches methane.

OBSERVATION – SPACECRAFT IMAGER

1. Presently over black blotches noted on Titan approach.

CSCC TO SPACECRAFT CENTRAL IMAGE PROCESSOR:

Task 8: Initiate Task 3 – Examine topography data over black areas.

Response: Spacecraft Central Image Processor.

1. Contour map shows 3 circular black areas to be conic depressions 10 km, 5 km, and 6 km, respectively, in depth.

2. Spectral data indicates third IR band reflectance to be 1.6 times surrounding average in these areas.

3. Thermal IR indicates 100° K at centers of cones decreasing to 50° K at edges.

HYPOTHESIS: CSCC – Assume black areas to be craters.

CSCC TO SPACECRAFT CENTRAL IMAGE PROCESSOR:

Task 9: a. Analyze area surrounding black regions for height.

b. Examine geological data for lineaments in surrounding area.

Response: Spacecraft Central Image Processor.

1. In annuli ($R_1 < r < 1.1 R_1$; $R_2 < R < 1.2 R_2$; and $R_3 < r < 1.15 R_3$) surrounding each black area (B_1 , B_2 , and B_3); depth decreases from inside edge to outside edge by 1 km, 0.5 km, and 2 km, respectively.

2. Radial lineaments exist inside annuli. None outside.

COMMENT: CSCC – Crater hypothesis reinforced, since very warm in center, assume to be volcanic in origin.

CSCC TO SPACECRAFT IMAGER AND CENTRAL IMAGE PROCESSOR:

Task 10: Initiate Task 5 – Analyze geological activity at terminator.

Response: Spacecraft Imager.

1. Texture in grey levels along terminator indicate mountain size shadowing.

Response: Spacecraft Central Image Processor.

1. Contour map indicates mountain range generally parallel to terminator from 70° south latitude.

REPORT: CSCC – Present knowledge indicates mountains can cause sharp cutoff of light. Unless contradictory evidence, assume mountain range at terminator causes observed sharp cutoff of light at terminator.

July 29, 2010

REPORT: CSCC

8:25 am GMT – Initial Planet Model completed. Atmospheric sounder indicates planned descent to surface procedure possible. Beginning procedure for finding highest priority landing target.

CSCC TO SPACECRAFT CENTRAL IMAGE PROCESSOR:

Task 51: Correlate parameters for most desirable landing site with surface model.

Response: Spacecraft Central Image Processor.

1. Highest priority areas – Those indicating possible carbon-based life or structures made by intelligent beings. – None indicated by Planet Model.
2. Next highest priority – Areas indicating possible H₂O. No indication of present H₂O by surface model. Possible ancient riverbed 25°-27°N. latitude and 34°-38° W. longitude. Topology map indicates possible target 25.1° N. latitude and 36.8° W. longitude. Target is within lander range of stratified river bank-like structures.

July 30, 2010

REPORT: CSCC

9:45 pm GMT – Have loaded site model into Lander Central Image Processor. Task sequences programmed to deploy parachute at 100 km, begin wind determination and parachute descent path modification at 45 km. Parachute to be ejected at 2 km. Landing rockets and obstacle avoidance imaging system to be initiated at 1500 m.

10:31 pm GMT – Initiating Lander Descent.

Response: Lander Central Image Processor.

1. Have locked onto target area.
2. Estimated ground drift and surface wind indicates parachute descent direction modification of 265° and 3.1 km necessary to hit target.
3. Obstacle avoidance system activated.

Response: Lander Guidance Image Processor.

1. Obstacle at site. Shift 30°, 0.16 km.
2. Site clean under lander, okay for vertical descent to surface.

July 31, 2010

REPORT: CENTRAL LANDER COMPUTER (CLC)

1:38 am GMT – Lander site assessment procedure initiated. Lander Guidance Imaging System turned on. Response: Lander Central Image Processor.

1. Surface immediately surrounding lander mostly small rocks on relatively flat surface. Hill (slope < 30°) blocking view beginning 100 m away 248° to 0°. Surface of hill easily navigable. Stratified rock wall beyond 1 km 0° to 20°. Hill obscures wall beyond 0°.

CLC TO LANDER GUIDANCE IMAGE PROCESSOR.

Task 1: Initiate analysis to find safe path to climb to apex.

Response: Lander Guidance Image Processor.

1. Stereo depth and contour data added to site model.
2. Safe path calculated.
3. Initiating journey.
4. Apex of hill reached. Stereo depth and contour data being added to site model.
5. Safe path possible in forward direction.

OBSERVATION -- LANDER CENTRAL IMAGE PROCESSOR.

1. Rock formation indicating upheaval at 240°.
2. No major obstacles indicated on Planet Model.
3. No major obstacles indicated in image pointed at 240°

CLC TO LANDER GUIDANCE IMAGE PROCESSOR.

Task 2: Initiate analysis to find safe path to formation.

Response: Lander Guidance Image Processor.

1. Safe path calculated for initial 100 m.
2. Beginning journey.
3. Dead reckoning from surface model and relative size indicates 0.5 of total distance covered.
4. Dead reckoning and size of upheaval in image indicates 300 m from upheaval.
5. Slowing down.
6. At base of upheaval. Rubble makes further progress in this path impossible.

CLC TO LANDER GUIDANCE IMAGE PROCESSOR.

Task 3: Initiate experiment No. 4379 – Rock specific density.

Response: Lander Guidance Image Processor.

1. Reconfiguring to manipulator vision configuration.
2. Located oval shaped rock 3 cm × 8 cm, not imbedded, within reach of manipulators.
3. Surface model for 0.5 of rock recorded.

4. Initiating manipulator to lift and weigh rock.

5. Rock weighs 15 N, Mass = 10 kg.

6. Initiating rotation of rock 180° with respect to initial position.

7. Surface model for remaining 0.5 of rock recorded.

8. Volume of rock is 0.010 m³. Density is 1000 kg/m³.

APPENDIX 3C

ILLUSTRATIVE HYPOTHESIS FORMATION SCENARIO

The scenario presented in table 3.6 suggests the great complexity of data processing and hypothesis generation involved in solving problems in a planetary investigation conducted by a fully-autonomous spacecraft. Table 3.6 shows a simulated report based on studies of the Martian oases following the Viking mission to the Red Planet in 1976 (Huguenin, 1978). Without delineating all logical

functions required to arrive at the final suggested hypothesis, the series gives the reader a feel for the many steps involved in full-fledged scientific analysis of a new situation. It is presented in the format of a condensed message sent to Earth via statements which indicate what measurements were made and confirmed, and what reasoning was used to draw specific conclusions.

TABLE 3.6.— HYPOTHESIS FORMATION: AN ILLUSTRATIVE SCENARIO

“Condensates appear suddenly at dawn in two different locations: $(-25^{\circ}, 85^{\circ})$ and $(-30^{\circ}, 315^{\circ})$.”

“Condensates act to flatten the reflectance spectra and appear as highly reflecting at blue wavelengths.”

“Blue cloud activity occurs during southern fall and winter, and mixed blue and yellow cloud activity occurs during spring and summer. Yellow clouds are indicative of dust. (This information will turn out not to be included in final hypothesis.)”

“In these two locations, condensates are brightest at dawn, indicating low-level hazes or frosts. If the brightening of the condensates occurred in the afternoon, convective cloud activity would be indicated.”

“Both areas are also major centers of dust storms, the clouds typically appearing suddenly at dawn and fading from white to yellow by noon.”

“Local winds are not sufficient to make airborne particles of the size observed at these sites.”

“There appears to be a correlation between the time and the location of appearance of both the condensate hazes and the dust storms.”

“Since winds cannot produce the dust storms, the hypothesis is that rapid evaporation just after sunrise of any water present in the soil is explosively ejecting dust particles into the surface atmosphere. The explosive action is the result of low ambient surface pressures and atmospheric densities.”

3C.1 References

Huguenin, R. L.; *et al.*: Mars: Infrared Spectral Reflectance and Compositional Implications. *J. Geophysical Research*, vol. 83, no. B11, Nov. 1978, pp. 5433-5441.