CHAPTER 6

TECHNOLOGY ASSESSMENT OF ADVANCED AUTOMATION FOR SPACE MISSIONS

A principal goal of the summer study was to identify advanced automation technology needs for mission capabilities representative of desired NASA programs in the 2000–2010 time period. Six general classes of technology requirements derived during the mission definition phase of the study were identified as having maximum importance and urgency, including autonomous "world model" based information systems, learning and hypothesis formation, natural language and other man-machine communication, space manufacturing, teleoperators and robot systems, and computer science and technology.

The general classes of requirements were individually assessed by attempting to answer the following sequence of questions in each case:

(1) What is the current state of the relevant technology?

(2) What are the specific technological goals to be achieved?

(3) What developments are needed to achieve these goals?

After the mission definition phase was completed, summer study personnel were reorganized into formal technology assessment teams with assignments based on interest and expertise. The results of this activity are summarized below.

6.1 Autonomous World Model Based Information Systems

The first assessment team considered the technology necessary to autonomously map, manage, and re-instruct a world-model-based information system, a part of which is operating in space. This problem encompasses technology needs for a wide range of complex, computerized data systems that will be available twenty or thirty years hence. The concept of a world model aboard a satellite operating without human intervention appears useful for a variety of satellite missions, but is specifically required for the terrestrial applications IESIS (Intelligent Earth-Sensing Information System, see chapter 2) and Titan exploration (see chapter 3) missions defined during the summer study.

The world model in space serves as a template by which to process sensor data into compact information of specific utility on Earth. It can consist of mapping data and modeling equations to describe, by past experience, the expected features the spacecraft will encounter. The use of the model requires algorithms in conjunction with the spacecraft sensors. A companion central model of higher sophistication will be required to further process, analyze, and dis-

seminate the information and to update the entire world model. In the IESIS this component is on Earth and in the Titan mission it is centered around Titan, but in either case the entire system requires autonomous management.

The following are some very general requirements for an autonomous world model-based information system; however, only the first two are discussed here.

- Development of mission specific philosophy for handling the mission data
- Model of the user and user requirements
- Realistic mission simulation techniques to test mission designs
- Modular satellite components
- Satellite serviceable in space
- Fault-tolerant design
- Autonomous navigation assistance
- Communications network
- Autonomous pointing, navigation, and control
- Standardized software to run and maintain satellites
- Data return

It is obvious that each NASA space mission should have specific information goals and that the data handling required in each must suit those goals. Costly data transmission and storage beyond that strictly required for mission operations should be eliminated. The sensor set adopted for a mission and its use must directly serve mission goals.

The goals of the Titan mission differ widely from those of the intelligent Earth-sensing system. In comparison with Earth, Titan is basically unknown. The space exploration mission goal is generally to explore and to send back as much general information as possible to terrestrial researchers about Titan. The Earth is much better known, so a major IESIS goal is to return very specific information in response to user requests or system demands. (In this latter mission, raw pixel data should be returned only under very restricted circumstances. Users requiring raw data should pay a premium for it and should accept archiving responsibility as well.) Each mission will develop a uniquely relevant data-handling philosophy. This, of course, presupposes that models are available of mission users who are the final recipients of the data.

TABLE 6.1.– SOME SIGNIFICANT LANDMARKS
IN WORLD MODEL CONSTRUCTION

Year	Landmarks
1988	Autonomous on ground construction and test of a world model directly from advanced Landsat data
1990	Shuttle demonstrations of intelligent satellite system begin
1990	Primitive world model for Titan mission
1992	Completed user models by opening advanced Landsat ground test to selected user
1994	Autonomous satellite demonstration
1995	Titan intelligent demonstration mission launch
2000	Start of Intelligent Earth Sensing Information System

Table 6.1 lists a few milestones in the production of a completely autonomous and sophisticated satellite world model system. The Titan mission proposed in this report would be scheduled for launch in about 1995 and the Earth-sensing system would go into operation in 2000 AD, although a more primitive version of the world model could be ready by 1990. Since Titan is largely unknown, its world-model system must be capable of constructing a database almost entirely from first-hand on-orbit observations of the planet, hence should most properly be termed a "modeler." The Titan modeler and Earth model initially will be developed autonomously on the ground using incoming imaging data from an advanced Landsat-type satellite using conventional computers, memory, and Space Shuttle demonstrations (Spann, 1980). Test operations will characterize the operation of world model systems, and as testing continues the Earth model portion can be opened to selected users for terrestrial applications purposes. User access will allow development of worthwhile user models for the forthcoming IESIS mission (Rich, 1979). If the world model programs are successful, launch of the Titan modeler could take place in 1995 and initiation of IESIS could begin in the year 2000.

The important features in the operation of the world model arranged from its internal database through its construction, sensing, management, and user interface are:

- Techniques for autonomous management of an Intelligent Satellite System

- Mapping and modeling criteria for creation of a compact world model

- Autonomous mapping from orbital imagery

- Efficient rapid image processing techniques against world models

338

- Advanced pattern recognition, signature analysis algorithms for multisensory data-knowledge fusion
- Models of the users
- Fast high density computers suitable for space environment.

Autonomous hypothesis formation and natural language interfaces are important additional techniques discussed in detail in the remainder of this report, and a summary of specific recommendations of the remaining sections are in the following categories:

1. Land and ocean models
2. Earth atmosphere modeling
3. Planetary modeling
4. Data storage in space
5. Automatic mapping
6. Image processing via world model
7. Smart sensors
8. Information extraction techniques
9. Active scanning
10. Global management of complex information
11. Systems plan formation and scheduling

6.1.1 Land and Ocean Database

Each world model is specific for a given mission goal. For a land-sensing Earth mission the satellite model may be as simple as a flat map with discrete "niches" specified by type, coordinates, rough boundaries, and nominal sensor and characteristic values. The niche type may be separately catalogued and a file stored of important niche characteristics, sensor combinations useful in determining boundaries between two niches, normal anomalies, and information extraction and sensor-use algorithms. Sensor combinations most useful in determining niche boundaries must be developed. The ground component of the model will be more advanced, combining finer detail, historical data, local names, seasonal and temporal information, and complex modeling equations. Oceanic (and atmospheric) components of the world model will require sophisticated dynamic representation.

The satellite model is the component of the world model used for direct on-board processing. Without the satellite component, it is not possible to accomplish the very large data reduction inherent in model-based systems. The satellite model must be stored so that it is compact, sensor specific, capable of updating, and consistent with its use in image processing and in the particular orbit overpass.

In the image processing on board the satellite, the large number of pixel elements spanning a niche in each sensor is replaced by a small set of niche sensor characteristics such as area, average value, variance, slope, texture, etc. A highly convergent representation of desirable descriptors is required so that these few niche-dependent characteristics can faithfully represent the multitude of pixel points.

Oceanic niches and cells are commonly quite large. The ground model is necessary as a master for updating and against which the various satellite models may be calibrated. The ground model includes the library and archiving functions of the whole system.

Land and ocean database technology requirements are:

- Identification and characterization of important niches on land, ocean, atmosphere, and in boundary regions between

- Optimum niche size for use in space image processing

- Determination of well separated, easily identified niches to serve as geographical footprints

- Compact representation of niche boundaries

- Optimum sensor combinations for each niche

- Optimum sensor combinations for boundaries

- Anomaly specifications for niches

- Convergent set of niche specific characteristics

- Nominal values for niche characteristics in each niche in each sensor and for various sensor combinations such as sensor ratios

- Dynamic models for temporal variations of land, ocean, and atmospheric niches

- Optimum coordinate system for storing world model for computer readout in strips during orbital pass

- Optimum distribution of a complex world model within a multicomponent system

- Advanced data cataloguing

- Models of the users and their requirements.

6.1.2. Earth Atmospheric Modeling

The choice of sensor measurements most appropriate for terrestrial meteorological monitoring will require great advances in our present understanding of the atmosphere. Because of their dynamic and highly interactive character, the boundaries of homogeneous atmospheric three-dimensional niches will be far more difficult to define than surface niches whose features are essentially stationary by comparison. An important stage in the development of the Intelligent Earth-Sensing Information System will thus be the definition of useful niche concepts. Choosing measurements important for monitoring the Earth's atmosphere will also require great advances in present understanding of both lower and upper atmospheric phenomena.

Lower atmosphere. Examples of possible lower atmosphere niches might be regions where (a) certain temperature or pressure regimes such as low-pressure cyclones are operative, (b) there is a concentration of a particular molecular species, or (c) there is a characteristic cloud pattern

indicative of an identifiable dynamic process. Such niches will often overlap, being highly interactive and transient. If the concept of a niche is to be efficient its boundaries should be essentially independent of its major properties, although property dependent niches could also be very useful. Lower atmospheric niches are time-varying in size and location, constantly appearing, disappearing, and merging.

The size of each niche will depend partly on the complexity of the atmospheric region and partly on requirements for effective monitoring or modeling. For example, atmospheric niches near the Earth's surface will undoubtedly be smaller than those in the upper atmosphere because of the complexity of surface weather patterns and of the need to have detailed niche descriptions to develop adequate meteorological models.

Properties measured in lower atmospheric niches will include a wide range of parameters — pressure, temperature, humidity, cloud cover, wind speed, rainfall, atmospheric components, etc. Each property will have its measured values processed within the three-dimensional niche in a useful niche-dependent manner. This may be used to extract data showing, for example, the average rainfall in a niche area, its gradient toward niche boundaries, patchiness of the rainfall pattern, and higher-order characteristics. Since the atmospheric niche sizes are large, the savings from averaging three-dimensional data can be huge. To ensure that niche properties such as rainfall are faithfully reproduced over the niche, the number of higher-order characteristics such as Fourier components of the data may be large, perhaps several hundred.

Another alteration of the land sensing concept must be made when comparing incoming observations to a resident world model on board the satellite. To meet Earth's needs, satellite descriptions of local weather should be continually updated together with models of the processes involved so that predictions may be made. The ephemeral and interrelated nature of many of the weather pattern-defining niches will make comparisons of current with previous observations difficult to interpret. The changing values and spatial extent of niches characterizing temperature, moisture, or pressure must be understood within the context of a complete model of weather activity. If local weather models are part of the resident world model, elaborate adaptive modeling must occur on board to correlate the incoming niche observations and to fit them into a model. In the case of lower atmospheric weather, it may be most efficient to transmit complete niche descriptions from every pass of the satellite for on-ground modeling to determine, say, the appearance of storms (high and low pressure areas) using complex pattern recognition algorithms, weather expert systems, and large computer storage.

Niches which are large or do not involve complex interfacing or modeling in conjunction with other niches lend themselves more easily to comparison with world land models. Changes in large-scale gradients and global trends in

temperature, pressure, particulates, cloud cover, and rainfall could be detected by continuous matchings with an atmospheric world model. On a smaller scale the detection of an increase in concentration of a particular molecular species (e.g., an SO_x pollutant) could also be made by simple comparison. Table 6.2 summarizes possible categories of large-scale spatial niches.

Upper atmosphere. Earth's upper atmosphere involves complex chemistry, photochemistry and transport processes. Although significant progress has been made in understanding these processes, there is still a great deal of uncertainty in present knowledge of the stratosphere, mesosphere, and lower thermosphere.

The upper atmosphere covers the range of 15 to 150 km in altitude. Absorption and emission of radiation occur over a wide range of the electromagnetic spectrum at these heights. Satellite systems presently exist or are in the planning stages which perform high resolution passive radiometry measurements in both down-looking and limb-sounding modes of vibrational (IR) and rotational (mm) molecular transitions. Limb-sounding microwave techniques will for the first time allow study on a continuous global scale since spectral lines are observed in emission. Micro-

TABLE 6.2.– POSSIBLE CATEGORIES OF PROPERTIES OF LARGE SCALE SPATIAL NICHES USEFUL FOR EARTH MONITORING

- Humidity profiles
- Precipitation location and rates
- Air pressure profiles and gradients
- Air temperature profiles and gradients
- Clouds – cloud top temperature, thickness, height, extent/location, albedo
- Atmospheric electrical parameters – lightning, magnetospheric electric field
- Atmospheric winds
- Aerosol size and concentration
- Particulate size and concentration
- Oxidant levels
- Molecular species – natural and man-generated

$CFCl_3$	HCl
CF_2Cl_2	HF
CF_3Cl	HNO_3
CH_4	H_2O
$ClONO_2$	HN_3
CO	NO
C_2O	N_2O
CO_2	SO_2
	O_3

- Also: Atmospheric transmittance, solar constant, solar flare activity, solar particle detection, Earth radiation budget

wave receiver technology is rapidly advancing to submillimeter wavelengths which will enable the measurement of many additional minor atmospheric constituents that play a part in radiative transfer processes. Distribution of such constituents is determined by various chemical and photochemical reactions and by atmospheric motions on both small and large scales.

The current research interest in the field of atmospheric studies reflects the present level of understanding of the atmosphere. Of particular importance are measurements improving the knowledge of how man's increasing technological activities may perturb stratospheric processes and affect the maintenance of the stratospheric ultraviolet-shielding ozone layer. These upper atmospheric studies require long-term precision composition and thermal measurements. An understanding of the role of the stratosphere in climatic change and atmospheric evolution is also needed. This includes understanding stratospheric warnings, their impact on chemistry, the spatial distribution of aerosols, and interactions with the troposphere below and the mesosphere above. Measurements of the mesosphere and lower thermosphere are needed to determine composition and variability. Little is known of the basic meteorology in these regions (temperature, pressure, wind variations). Possible variations in O_2 in these levels may affect ozone concentration at lower altitudes.

The long-term goal is the development of an intelligent Earth-sensing information system which can compare synopses of complex numerical models of the upper atmosphere with specific observations which are a subset to the original observations required to design those models. Comparisons could be simply the matching of predicted or acceptable values with observations. The actual models to be flown will have varying degrees of complexity. Most models may just be listings of predicted values derived from complex numerical models. These listings could be compared with observed values (for developmental purposes). Subsequently, measurements might be reduced to those spectral lines which yield the most information with optimum redundancy. For the purpose of testing systems which will be flown on the Titan mission it will be necessary to fly models which are or can be self-modifying to account for any observed discrepancies.

Since the Earth's atmospheric modeling will be done in much greater detail than is necessary for planetary exploration, tests of adaptive radiative transfer and hydrostatic equilibrium modeling should be kept simple. In planetary exploration, relatively crude remote sensing to determine composition, winds, atmospheric structure, cloud cover, and temperature profiles will be needed to obtain a general understanding of the planet's atmosphere. However, it may be valuable to include complex modeling systems to explicate possible organic chemistries.

To reach the required level of understanding of the atmosphere, extensive studies must be undertaken to

develop and validate complex models complete in their inclusion of aerial chemistry, distribution of minor constituents, radiation fields, and large-scale dynamics as a three-dimensional time-dependent problem. When the upper atmosphere is sufficiently understood, appropriate parameters to be monitored and modeled can be determined. Useful techniques for verifying models will involve checking model predictions with the observed distribution and concentration of chemically active species (some of which may also be useful as tracers of atmospheric motions).

Current planning for the versatile microwave limb-sounders seems to be moving in a direction compatible with Earth and planetary sensing requirements. The radiometers will be modularly constructed so that they can be easily exchanged as measurement priorities change and technology advances. Limb-sounder instruments will probably be capable of accommodating several radiometers for simultaneous measurements. Instruments in different spectral ranges will be employed for complementary measurements. The antenna, scanning, data handling, and power supplies should be common to any complement of radiometers used in the system.

The Earth atmospheric modeling technology requirements are:

- Definition of lower and upper atmosphere niches (spatial location or characteristic properties)

- Adaptive modeling of weather
  - complex pattern recognition algorithms
  - weather expert system

- Sensors for measuring lower atmospheric properties

- Determination of set of properties in an atmospheric niche which give consistent boundaries

- An understanding of the atmosphere sufficient to know what parameters need to be monitored
  - development of high resolution satellite microwave techniques for measurement of minor constituents

- Use of microwave limb-sounding techniques for continuous global coverage

- Development of an optimum sensor set for monitoring the upper atmosphere.

6.1.3. Planetary Modeling

For a relatively unknown body, surface and atmospheric modeling must evolve in greater detail during the course of the mission as more information on important planetary characteristics is obtained. A systematic methodology is required for understanding and exploring a new environment using high sensor technology. This methodology must determine what questions should be asked, and in what

order, to efficiently and unambiguously model an uncharted atmosphere and planetary surface. A decision must be made early in the planetary mission whether to place emphasis on elaborate remote sensing from orbit, which may ensure survivability but will not allow all of the scientific objectives to be met, or to physically probe the atmosphere, thus exposing a mission component to increased danger but allowing more precise determination of useful atmospheric properties.

The planetary probe must be capable of orbiting, investigating, and landing during a single mission. This is a difficult task to accomplish in one fixed design because of the uncertainties in the nature of the unknown planetary environment. The resulting planetary modeling requirements are:

- Systematic methodology for exploring an initially unknown environment

- Decision ability in the face of lethal dangers

- Modeling ability to establish norms of a planetary surface which allow recognition of interesting sites

- Autonomous creation and updating of planetary models using a variety of complementary sensors

- Adaptive programming of atmospheric modeling to establish atmospheric parameters

- Complex modeling or organic chemistry processes

- Expert systems for spectral line identification of complex and ambiguous species

- Develop general spacecraft capable of adaptation under uncertain atmospheric and surface conditions and which possesses a broad set of sensors

- Exchangeable radiometers, each capable of simultaneous measurements, and with wide spectral range and self-tuning ability

- Mass spectrometers and radio spectrometers based on range of organic compounds considered important or highly probable

- Instruments with interchangeable and reconfigurable basic elements

- Development of space qualified subsystems, instruments with long lifetimes

- Development of smart probe sensors and high speed image processors able to operate in the short period of time available during descent

- Use of sensors which record only significant variations in incoming data

- Simple redundancy so spacecraft will not be overloaded with back-up instrumentation

- Automated failure analysis systems, self-repairing techniques.

These requirements will engage numerous disciplines and thus create challenging instrumentation and design engineering problems for mission planners.

6.1.4 Data Storage

The terrestrial world model will require satellite storage of from 10^{10} to as much as 5×10^{11} bits and perhaps 10^{14} bits on the ground. Forecasts (Whitney, 1976) give estimates of 10^{14} bits of in-space memory and $>10^{16}$ bits of a typical on-ground memory by the year 2000. The data storage should be structured in a manner compatible with build-up of an image and extraction of image processing during orbital overpass. Optical disc, electron beam, and bubble memories are possible candidates in addition to more conventional alterable memories.

The technology requirements include a high density, erasable memory suitable for use in the space environment, optimum memory architecture for readout of the world model during orbit overpass, and error-correcting memory design.

6.1.5 Automatic Mapping

Terrestrial automatic mapping by IESIS can be accomplished using geographical data already obtained from Earth or from satellite data alone. The Defense Mapping Agency has developed digital mapping techniques for regions of the globe (Williams, 1980). By contrast, the mapping of Titan must be accomplished almost exclusively from orbit. In either case, information in the form of niche identification, basic modeling equations, and known planetary parameters will be supplied from Earth both initially and during operations. Automatic mapping from space requires state-of-the-art AI techniques including boundary and shape determination, optimum sensor choice, niche identification, and learning techniques. Full autonomous learning by abduction and inference to build new knowledge is presently beyond the capability of AI (see section 6.2). Though use of such advanced AI techniques would tremendously enhance the utility of a satellite world-model-based information system, they are not considered essential in this application.

Mapping technology ultimately must prove sensor-independent since the map produced should reflect a reality existing in the absence of the sensor data. However, specific sensor combinations will produce a completed map more rapidly and reliably depending upon the niche environment to be mapped. Orbits which repeat over fixed portions of the planet are especially advantageous in assisting automatic mapping and memory structuring.

Technology requirements, summarized briefly, are:

- Rapid autonomous mapping techniques from orbital data

- Optimum sensor combinations for reliable and rapid mapping
- Determination of relative advantage of radar, optical, IR mapping
- Optimum orbit height and orbit type for automatic mapping
- Techniques to rapidly, reliably, and automatically update world model components in satellites and on ground directly from orbital image data
- Digital mapping techniques
- Autonomous hypothesis formation techniques.

6.1.6 Image Processing via World Model

The satellite memory component of the world model is used for image processing. The actual image data from one or several sensors must be cross-correlated with a pass map (retrieved from memory) in strips along the orbit to produce an optimum match of imaged niches with their mapped locations. This process rectifies the sensed image and produces geometrical corrections necessary to adjust the sensed image to the reality reflected in the stored map. This process also will help determine the precise satellite location (Kalush, 1980). Boundaries must then be identified from the actual imagery and compared to the nominal boundaries of the map. The boundary area is an important and simply determined characteristic of the niche. Other characteristics such as anomalies are determined by new boundaries, altered location of boundaries, or changes in the determined sensor readings from their expected values. Temporal, sensorial, and solar corrections must be applied to the sensor readings and defining labels supplied for all niche characteristics for complete referencing purposes (Schlienn, 1979).

The satellite location can be combined with velocity and navigation information from a global positioning system to prepare for the next image in the sequence. This preparation allows minimum processing in the subsequent image rectification and permits determination of the optimum sensor combination for the next imagery. Instructions from Earth ground control or a central satellite autonomous manager must be incorporated into the preparation and image processing procedures.

Very sophisticated computer technology is required aboard the satellite to accomplish the image processing. Such processing is not found on any present-day satellites, and is done on-ground only in very limited form today. Fully parallel processing techniques are anticipated as a possible alternative to serial processing (Gilmore et al., 1979; Matsushima et al., 1979; Schappell, 1980). Optical processing methods should also be investigated since these

342

techniques are naturally parallel (Vahey, 1979). Technical advances are needed in the following:

- Automatic techniques to rapidly correlate memory stored mapping and modeling information with visual and radar imagery obtained in orbital pass

- Fast image enhancement and threshold techniques

- Rapid cross-correlation techniques

- Rapid boundary-determination techniques

- Rapid Fourier transform techniques

- Algorithms for improved automated data associations

- High density rapid computers for use in space environment

- Parallel processing computer techniques involving large wafers, advanced cooling techniques, advanced interconnection techniques between array elements, more logic functions between elements performed in one clock cycle and advanced direct data output from array to a central controller

- Ability to load and unload imaging data in full, parallel manner at all stages of handling raw data

- Investigation of possible use of optical processing techniques such as holographic process or integrated optics for satellites processing of imagery via world model

- Techniques to rapidly, reliably and automatically update world model in satellite and on ground directly from image data

- Advanced data compression and compaction techniques for data transmission and storage.

6.1.7 Smart Sensors

Complex sensor configurations are required for both IESIS and Titan missions. A high degree of autonomous sensing capability is needed within the sensors themselves (Haye, 1979; Murphy and Jarman, 1979). These sensors must be smart enough to perform automatic calibrations, compensations, and to reconfigure themselves automatically — tasks requiring advanced memory capabilities and operating algorithms (Schappell, 1980). Desirable characteristics of such smart sensors on satellites (Breckenridge and Husson, 1979) are:

- Introduces no anomalies into data

- Performs analytical and statistical calculations

- Performs all operation in simplest form

- Adapts (handles) new data acquisition and processing situations

- Interactive sensor configurations

- Adjusts to different environmental conditions

- Extracts information in a useful form

- Makes decisions.

The use of a world model in conjunction with smart sensors would confer an extraordinary degree of intelligence and initiative to the system. In order to mate sensors most efficiently with the world model, the model should itself possess models of the sensor components. Since the sunlight at Titan is weak and the planet cold, efficient, visible, and infrared sensors are also necessary.

Technology requirements of smart sensors are:

- Advanced efficient solid-state imaging devices and arrays

- Sensor operation at ambient spacecraft temperature

- Electronically tunable optical and IR filters

- Advanced automatic calibration and correction techniques

- Distributed processing sensors

- Rapid, high responsivity detectors in near IR up to 3 μm

- Optimum set of sensors arrays for particular planetary mission

- Sensor models

- Silicon-based sensors with dedicated microprocessors and on-chip processing

- Investigation of piezoelectric technology for surface wave acoustic devices

- Sensor sequence control which can adapt to conditions encountered

- Precision pointing and tracking sensor mounts.

6.1.8 Information Extraction Techniques

Information can be extracted from sensory data originating from an object by recognizing discriminating features of the object. Such features are of three kinds: (1) physical (color), (2) structural (texture and geometrical properties), and (3) mathematical (statistical means, variance, slope, and correlation coefficients).

Humans generally use physical and structural features in pattern recognition because they can easily be discerned by human eyes and other common means. However, human sensory organs are difficult to imitate with artificial devices so these methods are not always effective for machine recognition of objects. But by using carefully designed algorithms, machines can easily extract mathematical features of patterns which humans may have great difficulty in detecting.

The algorithms will often involve a fusion of knowledge across multisensor data. As an example, the radiance observed from an object is a function of its reflectance,

343

incident illumination, and the media through which it is viewed. The ratio of the radiances at two different wavelengths can be used to separate water and vegetation from clouds, snow, and bare lands (Schappell and Tietz, 1977; Thorley and Robinove, 1979). (The radiance ratio for clouds, snow, and bare land is essentially the same, so these features must be separated on the basis of absolute radiance.) These two sensor procedures can be improved by using data from several sensor bands simultaneously in a multidimensional sensor space. Machines can process such complicated algorithms and "see" clusters in higher dimensions. The higher the multidimensional volume, the more accurate the discrimination between closely related sensorial characteristics.

The intelligent use of a world model requires autonomous, real-time identification of niches (through their features) and determination of characteristics. Real-time pattern recognition and signature analysis also must be accomplished to supply useful information to the user. Algorithms should be developed for identification, pattern recognition, and signature analysis.

Statistical procedures arise naturally in various classification schemes because of the randomness of data generation in various pattern classes. Statistical theory can be used to derive a classification rule which is optimal because it yields the lowest probability of classification error, on average. Various studies have developed decision functions from sets of finite sample patterns of classes. These decision functions partition the measurement space into regions containing clusters of the sample pattern points belonging to one clan. Some clustering transformations have been used in the development of such functions. Once a function has been selected, the main problem is the determination or estimation of its coefficients. For efficient coefficient estimation, time-dependent training samples are needed.

A wide variety of additional algorithmic techniques are needed. For example, texture analysis can be accomplished using gray-tone statistics and the time rate of change of spatial contrast along scan lines to distinguish among wheat, rye, and oats (Haralick et al., 1974; Mitchell et al., 1977). Below is a summary of technology requirements:

- Rapid methods for area centroid and orientation determination
- Rapid partitioning of image features
- Motion and relative motion detection
- Development of wide range of classification algorithms for user-defined applications
- Multispectral signature ratioing analysis and multisensor correlations
- Rapid texture analysis
- Investigation of usefulness of focal plain transformations for satellite use

- Schemes to allow disparate algorithmic techniques to interact to speed recognition process
- Determination of parameters of decision functions for various classification schemes

6.1.9 Active Scanning

The sensors discussed to date have been essentially passive — they do not generate the radiation they detect. For a variety of purposes, some satellite systems will engage in active scanning by highly efficient RADAR or LIDAR, all-weather imagery, night-time imagery, absolute and differential height determination, absolute and differential velocity determination, atmospheric probing, and leading-edge scanning. Of course the mission to Titan, relatively far from the Sun, will not have large amounts of power available for this purpose.

Additional technology requirements include a fast, efficient computer for generating imagery from SAR, the ability to determine height differentials to within several centimeters at boundaries, and the ability to determine differential velocities to within about 1 km/hr at boundaries.

6.1.10 Global Management of Complex Information Systems

Each mission explored by the study group consists of a very large, complex array of equipment and people widely geographically distributed, all of whom must work in a cooperative and coordinated fashion to achieve mission objectives. An important concern thus becomes the overall architecture of such a system, the way decisions are made and communicated, the coordination of tasks within the system, and the flow of information. These types of difficulties are not new in human endeavors and have been addressed within several disciplines which focus on specific aspects of the problem. A brief review of relevant fields resulted in several recommendations for high priority research in systems theory and control, summarized below.

Classical control theory. Systems which evolve according to well behaved physical laws describable in the form of differential equations have long been the domain of classical control theorists. The aerospace industry has been a prime user of this technology in the guidance and control of missiles and in the development of automatic pilots for aircraft. The system is usually modeled as shown in figure 6.1 which envisions an idealization of a physical system subject to stochastic disturbances (typically Gaussian). The system is observed and digressions from the preferred trajectory are noted. A controller working with the idealized model (expressed in the form of differential equations) and a specific objective (such as "hit a target within a given
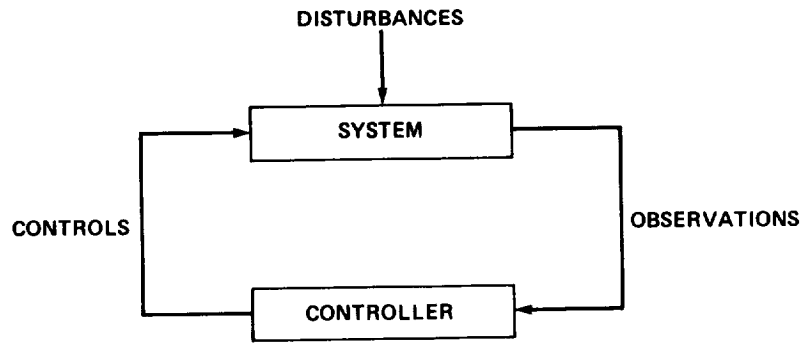
344

Figure 6.1.— Classical control theory systems model.

tolerance with a limited fuel budget") typically expressed mathematically in a quadratic form computes linear control to correct the system trajectory to meet the stated objective. This type of formulation is known as the LQG (Linear Quadratic Gaussian) formulation and has received wide attention within the control theory community.

Clearly, this theory is applicable to navigation and process control problems but will make only a rather minor contribution to the theory of how systems operate as a whole. This is not considered a critical mission technology since it is a fairly well developed and active field. Further, application depends on the notion of a single centralized controller. This is appropriate for micro control applications but inappropriate for macro control of large decentralized systems.

Game theory. Systems which employ multiple decisionmakers have been addressed by game theorists. Much of this work has been defense-related although economics has also provided an applications base. The basic notion of game theory is that there are an arbitrary number of decisionmakers, each of whom has an individual objective-function which may be (and likely is) in conflict with the objectives of the other decisionmakers. Each decisionmaker attempts to develop strategies which independently maximize the "payoff" to himself. Much work has been done on the "zero sum game" in which one decisionmaker's gain is another's loss.

If one envisions a cooperative, coordinated mission scenario, the current focus of game theory on threat strategies more appropriate to hostile environments seem ill-suited to peaceful space activities. A more appropriate meta-model is required for NASA's applications which reflects the necessity for cooperative coordination among the men and machines of the mission.

Nonclassical information control theory. The decentralized control problem for large-scale systems with a common (or at least coordinated) objective has received increasing attention in recent years. Initial work on "team theory"

(Radner, 1962) has centered on a team which is considered to have as its fundamental problem the coordination of decentralized activities utilizing delayed and imperfect information. The meta-model employed appears to be appropriate for the large-scale space missions considered in this report. Team theorists envision an autonomous ensemble of decisionmakers, each of which senses a local environment ("perfect" information) and can communicate in a delayed fashion with other decisionmakers ("imperfect" information). The ensemble, or team, shares a common objective and attempts to communicate as necessary for collective progress toward that objective to be optimized in some sense. This leads to the notion of an information structure among the members of the team.

The team concept has since been adopted within the control theory community and has led to "nonclassical control theory" — control theory which addresses multi-decisionmaker types of problems (Ho, 1980; Sandell et al., 1978). Much of this work is supported heavily by the Department of Defense (DOD) and focuses on problems of little direct relevance to NASA. Vigorous support by NASA of work in nonclassical control theory is recommended to develop more appropriate theories for the types of systems which comprise the space missions of the future. For instance, much of the DOD work addresses guidance and control problems, whereas NASA's prime interest would be more appropriately in information systems control. Supporting disciplines include probability and Markov decision processes. These are areas which are required to advance the state of the art in systems theory and control and to apply it effectively to NASA missions.

Prior work in these fields tends to focus on performance optimality as an objective. While optimality is a laudable goal, it is not clear that this should override other concerns such as stability and performance predictability. The formal tools currently available to evaluate the stability of a large decentralized system are virtually nonexistent. The major recommendation of the study group in this area is that NASA seriously consider the system-wide objectives of its future systems and support a program of basic and applied

345

research which develops the theoretical basis to achieve these objectives. The major relevant disciplines include, as a minimum set, nonclassical control theory, probability, queueing theory, and Markov decision theory. Technology requirements include the determination of system-wide objectives of missions and the development of theoretical and practical bases to achieve those objectives; development of nonclassical control theory of complex man-machine information systems; probability theory applicable to complex information systems; and Markov decision theory for complex information systems.

6.1.11 Plan Formation and Scheduling

In those cases where robots are called upon to operate in sophisticated task environments, the machine system first performs some computation which can be considered problem-solving, then takes action based upon the problem-solving result which is called the "plan formation" process. The part of the resulting plan which identifies the times at which actions are to occur is called the schedule. Whether the machine system is a relatively small mobile robot as might be used in planetary surface exploration, or a large distributed intelligence such as an Earth-sensing information system, several common features dominate in achieving effective and flexible operation (Sacerdoti, 1979):

- The ability to represent the state of the relevant parts of the world (the "world model")

- The deductive ability to recognize consequences of a particular world state description

- The ability to predict what changes will occur in the world state, possibly due to some action or actions a complex autonomous system itself might perform.

In most realistic environments it is impossible to completely build a detailed plan and execute it in unmodified form to obtain the desired result. Several difficulties preventing such a direct line of attack are:

(1) The external reality may not be known in sufficient detail to accurately predict the outcome of some action. If the action in question is the final one in a plan, then it may not achieve the overall intention of the plan. If it is an earlier action in a several-step plan, then it may not produce a required intermediate state for the overall sequence of actions to achieve the goal of the plan. If the goal is to make an observation to obtain information about the environment, the information obtained may not be adequate.

(2) Even if a perfect, or effectively perfect, model of the external environment is available to the robot, there may still be inaccuracy associated with the robot's control of itself (e.g., mechanical inaccuracy of motion).

(3) Other agents, with goals of their own, may alter the environment in unpredictable ways before the robot can complete the execution of its plan. In such cases some form

of overall coordination is necessary. It is not adequate simply to have the main goals of all of the active agents compatible. Even with this precaution, it is still possible to have a contention for resources or intermediate configurations in achieving the common goal. Aside from the problem of avoiding explicit conflict among several active agents there is the inverse problem of achieving efficiency increases by proper cooperative action among the agents.

For these reasons, a robot must continually monitor the results of its actions during plan execution, and modify the plan — in essence, re-plan — during plan execution.

A further complication arises when the plans must meet real-time constraints — that is, definite short-term requirements for actions where failure to meet the timing requirements carries significant undesirable consequences. Two types of real-time constraints, "hard" and "soft," may be distinguished. A "hard real-time constraint" is that the failure to carry out a successful plan that attains the relevant goal within the limits will result in a consequence so undesirable that extreme care must be taken not to overrun the time boundary. An example in the area of large-scale space construction might be the joining of two relatively massive but fragile substructures. Failure to initiate timely deceleration of substructures approaching each other could result in large economic losses. An example of a "soft real-time constraint" is in the maximization of the utilization of a costly resource, such as the observation satellites in an Earth-sensing system where it is important to schedule observations in such a way so as to minimize the number of satellites necessary to provide a given level of observational coverage. In this case, each individual failure to meet the real-time constraints has, in general, only minor consequences, but a continuing high-frequency of failure will result in economic losses through inefficient operation.

Because of the need to re-plan during plan execution, and because of the necessity to meet real-time constraints, it is important that complex autonomous systems have plan formation capabilities well in excess of current state of the art.

Current assessment. A considerable amount of work has been done in AI on problem-solving in general, and on planning and plan execution in particular. In the last 10 years the problem-solving emphasis has shifted away from planning towards the perceptual processes of vision and speech recognition. Table 6.3 lists some techniques for problem-solving and planning, and various representational schemes (NASA SP-387, 1976).

The frame notion of Minsky initially generated much interest and discussion, but little has been accomplished to date in terms of applications. There are attempts from several different perspectives to implement frame-based languages for programming, as for example KRL (Bobrow and Winograd, 1977), FRL (Goldstein and Roberts, 1977),

TABLE 6.3.— FORMS OF REPRESENTATION AND PROBLEM-SOLVING TECHNIQUES USED IN ARTIFICIAL INTELLIGENCE

Representations

- State space (Van de Brug and Minker, 1975)
- First order predicate logic (Nilsson, 1971)
- Semantic nets (Woods, 1975)
- Procedural embedding of knowledge (Hewitt, 1970)
- Frames (Minsky, 1975)
- Production rules (Newell, 1963)

Problem-solving techniques

- Backtrack programming (Goulomb and Baumert, 1965)
- Heuristic tree search (Pohl, 1977)
- GPS means-ends analysis (Newell, 1963; Ernst and Newell, 1969)
- Problem reduction (Amarel, 1968)
- Theorem-proving (Nilsson, 1971)
- Debugging almost correct plans (Sussman, 1975)
- Procedural (Hewitt, 1970)
- STRIPS (Fikes and Nilsson, 1971)
- ABSTRIPS (Sacerdoti, 1974)
- Cooperating knowledge sources (Erman and Lesser, 1975)
- Rule-based systems, expert systems (Shortliffe, 1976)

---

and MDS (Srinivasan, 1976). These attempts were ambitious, and while all met with serious difficulties, the possibility remains that the problems can be overcome. An ideal frame-based programming language could make it easier to structure knowledge into larger coherent units than would otherwise be practical.

The controversy between procedural and declarative philosophies of embedding knowledge has dwindled. It is now realized that each has its particular function to perform in an overall system, and that neither alone nor in combination are they an adequate underlying basis for AI theories or for sophisticated program organization.

There is a growing trend toward considering the first-order predicate calculus, or minor modifications of it, as the primary mode of representing declarative knowledge in AI systems. The reason is that this calculus has a well defined semantics and other declarative representation schemes tend to be simply different notional systems struggling to capture the same semantic notions as the predicate calculus.

Interest in formal theorem-proving techniques has remained high, and perhaps has even increased slightly, despite the slow progress in increasing the efficiency of mechanical deduction. While theoretical understanding of mechanical theorem-proving is increasing, to date there is little advancement in efficiency beyond that of a decade ago. Theoretical work on model use in theorem-proving has progressed only slightly (Reiter, 1972; Sandford, 1980), and applications methodologies are nonexistent.

Theoretical work has progressed in using first-order Horn logic as a programming language (Kowalski, 1974). Horn logic is a subset of the first-order predicate calculus in which a large number of interesting problems can be expressed. A truly unexpected development is the successful implementation of a workable programming system for Horn logic in which several nontrivial programs have been written (Warren and Pereira, 1977).

Much interest has developed in a rule-based type of knowledge embedding for restricted domains. These systems are commonly called expert systems, and have shown interesting and relatively strong problem-solving behavior. A variety of reasoning task domains have been implemented (Feigenbaum et al., 1971; Shortliffe, 1976), and the rule-based knowledge embedding method is robust in its performance. However, several severe defects of such systems must be addressed before realistic problem domains can be adequately handled. Defects include extremely limited domains of application, the large efforts required to construct the knowledge base, and the inability to access a basic theory and perform an *a priori* analysis. Work is in progress to devise systems avoiding these particular problems (Srinivasan and Sandford, 1980). There are also some relatively minor human interfacing problems with the present systems (see section 6.3).

There is a general increased awareness of the importance of the role of meta-knowledge (knowledge *about* knowledge) in problem-solving and in planning. The important related area of reasoning relative to open world databases is just beginning to be investigated (Reiter, 1980).

The general problem of representing the external world in an appropriate machine representation is a fundamental unsolved problem. While many facts are representable in many ways, no known representation is adequate to handle even such common phenomenon as a glass of water falling to the floor and breaking. It is likely that a fundamental shift in current approaches is required to achieve adequate representations for much of "common" world knowledge. There is little indication at this time what these new approaches should be. However, certain such "common" world knowledge is at least partially tractable with current techniques as, for example, the acquisition and use of knowledge about large-scale spaces (Kuipers, 1977).

Identification of critical research areas. Table 6.4 lists a set of critical research areas in the general AI fields of problem-solving, plan formation, scheduling, and plan

TABLE 6.4.– CRITICAL AI RESEARCH AREAS IN ROBOT PROBLEM-SOLVING AND PLANNING

1. General robot reasoning about actions
2. Combining AI problem-solving and plan formation with operations-research-scheduling techniques and real-time constraints
3. Techniques for classifying problems into categories and selecting the appropriate problem-solving method to apply to it
4. Expert systems
5. Generalized techniques for dynamic accumulation of problem-specific knowledge during a problem-solving attempt
6. Techniques for abstraction, and the use of abstraction for search guidance
7. Methods of combining several representations and search techniques together in a coherent manner
8. Ways to structure systems to have both fundamental theories to allow *a priori* reasoning along with a procedural level of skill to allow efficient real-time response
9. Models and representations of reality

execution. Table 6.5 gives the relevant mission requirements in these areas, the missions to which they might apply, and the identification of which items from table 6.4 are most relevant.

Recommended actions. Traditionally AI has been predominantly a research-oriented activity which implemented systems primarily for experimental purposes. There is a growing awareness among AI researchers that the time has come to produce limited capability but useful working systems.

In like manner, NASA should obtain experience at the earliest possible date with elementary space-robot systems in such areas as fully automatic spacecraft docking and sophisticated Earth-sensing satellites. Theoretical research in AI problem-solving and planning techniques will be an active area for several decades to come. If NASA is to become effective in directing this research toward its own goals, then early experience is necessary with elementary state-of-the-art techniques — although substantial advantages can be obtained even with relatively unsophisticated, near-term AI planning and execution monitoring techniques. Most of the areas listed in table 6.4 will progress both at the theoretical and applications levels without NASA taking action. This theory will generally be supportive of NASA's needs, particularly that done by DOD for space applications. Communication between NASA and DOD is thus important in overall planning for both organizations. While DOD interests in the mission requirements listed in table 6.5 are likely to be restricted to categories

TABLE 6.5.– CORRELATIONS BETWEEN MISSION REQUIREMENTS, MISSIONS AND RESEARCH AREAS FROM TABLE 6.4

Mission requirement, MR	Mission	Relevant areas from table 6.4
1. Automated housekeeping functions for long-duration spacecraft	TM,[a] ES[b]	1,2,4,5,7,8
2. Fully autonomous sequencing of observations, active and passive, from orbit, from landers, and during interplanetary flight, for a variety of sensors	TM,ES	2,3,4,7,8,9
3. Automatic docking, refueling, repair and maintenance of semi-independent probes	TM, ES	1,3,4,5,6,7,8,9
4. Automatic deployment of landers and orbiters from a central orbiter bus or busses	TM	1,4,9
5. Automatic landing capability on a planetary body where the lander is physically designed as a general-purpose lander capable of achieving planet fall on planets with a variety of atmosphere densities, wind velocities, and surface characteristics	TM	1,2,9
6. Automatic sample-taking of atmosphere and soil samples, and automatic low level sequencing of a variety of chemical and physical analysis techniques	TM	1,3,4,5,8

[a]Titan mission.
[b]Earth mission.

MR1, MR2, MR3, and possibly MR4, these cover most of the research areas from table 6.4. If this is indeed the situation with respect to DOD, then NASA can concentrate primarily on implementation projects.

However, certain needs and operating scenarios are peculiar to NASA and are not likely to develop in theory or applications without direct NASA guidance. Two very

specific examples are the development of robot patterning techniques and the development of "show and tell" robot control (see section 6.3). While commercial industrial robots have long employed patterning methods, these methods are used only in a rudimentary form and further applications technology development is needed for them to become useful to NASA. The show and tell mode of robot action has apparently not been identified and investigated to any large degree, and seems to be an area where NASA should take an immediate and large interest both in the theory and the applications aspects.

6.2 Learning and Hypothesis Formation

The Titan exploration mission description, documented in chapter 3, discusses the characteristics of a machine intelligence system possessing autonomous self-learning. This capability, its relation to state-of-the-art AI, and the new research directions it demands are summarized below.

6.2.1 Characteristics

For a machine to learn a previously unknown environment involves both the deployment of knowledge structures correct for known environments and the invention (or discovery) of new knowledge structures. A machine intelligence system which learns could formulate (1) hypotheses which apply existing concepts, laws, theories, generalizations, classification schemes, and principles to the events and processes of the new environment, and (2) hypotheses which state new concepts, laws, and theories whenever the existing ones are inadequate.

Different logical patterns of inference underlie the formation of these types of hypotheses. Analytic inferences support the formation of hypotheses which apply existing concepts, laws, and theories. Inductive and abductive inferences support the invention of hypotheses which state *new* concepts. Analytic, inductive, and abductive inference are mutually and logically distinct — one of them cannot be replaced by some combination of the others (see section 3.3.3 and compare Fann, 1970; Hanson, 1958; Lakatos, 1970a, 1970b, 1976; Peirce, 1960, 1966).

6.2.2 State-of-the-art in AI

State-of-the-art AI lacks adequate and complete treatments of all three inferential classes necessary for the development of machine intelligence systems able to learn in new environments.

Analytic inferences receive the most complete treatment. For instance, rule-based expert systems can apply detailed diagnostic classification schemes to data on events and processes in some given domain and produce appropriate identifications (Buchanan and Lederberg, 1971; Duda et al., 1978; Feigenbaum, 1977; Martin and Fateman,

1971; Pople, 1977; Shortliffe, 1976). An expert system such as PROSPECTOR can identify a restricted range of ore types and map the most likely boundaries of the deposit when given survey data about possible ore sites (Duda et al., 1978). However, these systems consist solely of complicated diagnostic rules describing the phenomena in some domain. They do not include models of the underlying physical processes. In general, state-of-the-art AI treatments of analytic inference fail to link the detailed classification schemes used in these inferences with the fundamental models required to deploy this detailed knowledge with maximal efficiency.

Inductive inferences receive a less complete treatment although some significant advances have been made. For example, Hajek and a group of co-workers at the Czechoslovak Academy of Sciences have, over the past 15 years, developed and implemented systems of mechanized inductive generalization (Hajek and Havranek, 1978). They do not take the approach of "inverse deduction" which has been explored by Morgan (1971, 1973). Instead, the Czech group has developed techniques for moving from data about a restricted number of members of a domain, to observation statement(s) which summarize the main features or trends of these data, to a theoretical statement which asserts that an abstractive feature or mathematical function holds for all members of the domain. (For instance, see table 6.6.) Though they allow a role for what they call "theoretical assumptions," in moving from observation to theoretical statements they have concentrated their work on formulating the rational inductive inference rules for bridging the gap between the two —

TABLE 6.6.— SAMPLE INFERENCE DATA

Rat no.	Weight, g	Weight of rat kidney, mg
1	362	1432
2	373	1601
3	376	1436
4	407	1633
5	411	2262
Observation statement	Therefore, the observed weights of the kidneys have the same order as the weights of the rats with one exception.	
Theoretical statement	Therefore, the weight of a rat's kidney is positively dependent on the weight of the rat.	

though it is not clear that their system captures the full range of influence which fundamental models exercise over inductive inference (see section 3.3.3). An independent research effort in the United States attempts to integrate fundamental models with specific abstractive, or generalizing, techniques (Srinivasan, 1980; Srinivasan and Sandford, 1980). However, unlike the Czech group the American team is still at the stage of theory development – a working system has yet to be implemented in hardware.

Abductive inferences have scarcely been touched by the AI community, but nevertheless some tentative first steps do exist. A few papers on "nonmonotonic logic" were delivered at the First Annual National Conference on Artificial Intelligence at Stanford University in August 1980 (e.g., Balzer, 1980), and much discussion followed. However, this attempt to deal with the invention of new or revised knowledge structures is hampered (and finally undermined) by their lack of a general theory of abductive inference – with one notable exception, the recent work of Frederick Hayes-Roth (1980). Hayes-Roth takes a theory of abductive inference developed by Lakatos (1976) for mathematical discovery and makes two of the low-level members of the family of abductive inferences which Lakatos identifies operational. Still, this work is only a preliminary step toward implemented systems of mechanized abductive inference and, unfortunately, it seems to represent the extent of theory-based AI work on abductive inference to date.

In summary, state-of-the-art AI treatments of analytic and inductive inference provide no fundamental models as a theoretical foundation to support the detailed knowledge structures and inference techniques upon which the treatments are built. Yet these models are an essential and integral element of analytic and inductive inferences. State-of-the-art AI virtually lacks treatments of abductive inference. However, model-based analytic and inductive inference systems and an abductive inference system are all necessary prerequisites for machine learning systems.

There appears to be a growing acceptance within the AI community of the above problems and that overcoming these gaps in current treatments of analytic, inductive, and abductive inference is an important future research direction for the entire field. For example, at the First National Conference on Artificial Intelligence, Peter Hart admitted that the fact that rule-based expert systems lack a fundamental model to ground the detailed rules makes them superficial and inflexible. Charles Rieger at the University of Maryland is beginning to address the question of layering models under rule-based systems. Several recent AI initiatives with respect to inductive and abductive inference have already been noted.

A concerted and serious attack on the problem of developing a theory of abductive inference for machine intelligence could pay enormous dividends. First, machine learn-

ing systems cannot possibly possess a full learning capability unless they can perform abductive inferences. Second, a successful mechanization of abductive inference would require the solution of problems which must also be solved for the successful mechanization of analytic and inductive inference. These problems include: (1) how to represent the fundamental models of the processes which underlie the detailed occurrences of domains, (2) how to inferentially relate these to more detailed knowledge structures such as laws, principles, generalizations, and classification schemes, and (3) how to map the representations of a domain occurrence in one "language," say, that of the model, onto its representation in another "language," say, that of a set of diagnostic rules. Since an investigation of abductive inferences seems to hold many keys to solving the problem of machine learning, and since recent developments in AI seem to promise receptivity to such an investigation, the development of a theory of abductive inference for machine intelligence appears to be the preferred research direction for work on machine learning systems.

6.2.3 Two Barriers to Machine Learning

Two points from the above discussion must be emphasized. First, state-of-the-art AI work on hypothesis formation is almost totally devoid of research on abductive inference. However, machine systems must have this capability in order to be true learning systems. Second, current AI work on analytic and inductive inference tends to proceed in the absence of relevant theories, and this seems to be the reason why state-of-the-art AI treatments fail to give fundamental models their proper role in inference systems. However, adequate theories of all three types of inference are a necessary foundation for successful machine learning systems.

Both of these barriers to machine learning – the abductive inference barrier and the theory barrier – must be bridged before machine intelligence systems can be given a full learning capability. The abductive inference barrier has already been fully treated, but some additional discussion of the "theory barrier" is useful here.

Historically, technology has developed in two distinct patterns – empirical and theoretical. Empirical technology is a "black box" approach. Given the problem of producing action A from some set of inputs $(I_1, ..., I_j)$, it leaves the real-world process connecting $(I_1, ..., I_j)$ with A unanalyzed. Because a theoretical model of the process is not available, rules for producing A must be obtained exclusively by empirical discovery. For instance, gunpowder was discovered and utilized by people who did not have a theory of combustion adequate to explain chemical explosive action. Various steelmaking technologies were developed by medieval European and Arabian smiths in the complete

absence of an understanding of how and why their techniques worked. However, given the same problems, theoretical technology utilizes a theoretical model of the real-world process connecting $(I_1, ..., I_j)$ with A and derives rules of production for A from the model. Examples of theoretical technology include radar, lasers, the Polaroid Land camera, digital computers, and integrated circuits.

Although these two patterns are distinct, many specific technologies have a "mixed mode" pattern of development. In such cases, a model of the full process connecting $(I_1, ..., I_j)$ with A is still not available, but refinements and extensions of the empirically discovered rules of production are based on partial models of the process. That is, some decomposition of the full process into its subprocesses has been made, and models for these subprocesses have been constructed. This is not true theoretical technology, however, because no general model of the full process is available and, consequently, an integrated set of model-derived rules of production is not possible.

Empirical technology, but not theoretical technology, is ultimately self-limiting within any given field of technology. That is, there is a level of technological capability beyond which empirical techniques cannot penetrate. This level is a function of empirical technology's pattern of development, *not* of the world itself. The reason for this self-limiting characteristic is the absence of theoretical models. Empirical methods develop via trial and error, small incremental refinements and extensions of empirically discovered rules of production. Since the rules are not based on a model of real-world processes, however, these modifications cannot be orchestrated and integrated, but are instead *ad hoc* "fixes" that hold only over a limited domain. Once the modified empirically based rules of production reach a sufficient level of complexity, the probability becomes very high that the next *ad hoc* "fix" may undo a previous one. Further development in the particular technological field (development in the sense of increased technical capability) stops at this point.

Theoretical technology need not be self-limiting. Since it is based on a model, the above effect may not be present. Theoretical technology is thus able to push technological development in a given field to the maximum extent consistent with whatever real-world limitations characterize the field.

This discussion sets the stage for a consideration of the type of intelligence capability which can realistically be expected from machine intelligence research. The question of machine intelligence has been replaced by the question of the machine formulation of hypotheses. If we define a scale of hypothesis formulating capability (HYP) as $HYP \equiv TH + CRED$, where TH is the theoretical content of the hypothesis and $CRED$ is the credibility of the hypotheses, then the design goal for advanced forms of machine intelligence is to be as high on this scale as possible.

Either an empirical technology or a theoretical technology pattern can be followed in developing machine intelligence. However, with respect to the HYP-level which can be achieved by the two different patterns, the empirical technology approach is ultimately self-limiting at a level of hypothesis formulating capability which is lower than that prerequisite for automated space exploration (see sections 3.2 and 3.3). It is clear that automated space exploration and other applications requiring very advanced machine learning are possible only if the theoretical technology approach to machine intelligence is employed.

Unfortunately, AI is currently taking the empirical technology approach to hypothesis formulation. There is nothing mysterious about the theoretical approach — it may be started by research into the 'patterns of logical inference by which hypotheses are formulated. Such an approach is limited only by the degree to which hypothesis formulation is logical and inferential. On the condition that it is, then the theoretical technology does not face a real-world barrier to achieving full machine-hypothesis generating capability.

6.2.4 Initial Directions for NASA

Several research tasks can be undertaken immediately by NASA which have the potential of contributing to the development of a fully automated hypothesis formulating ability needed for future space missions:

(1) Continue to develop the perspective and theoretical basis for machine intelligence which holds that (a) machine intelligence and especially machine learning rest on a capability for autonomous hypothesis formation, (b) three distinct patterns of inference underlie hypothesis formulation — analytic, inductive, and abductive inference, and (c) solving the problem of mechanizing abductive inference is the key to implementing successful machine learning systems. (This work should focus on abductive inference and begin laying the foundations for a theory of abductive inference in machine intelligence applications.)

(2) Draw upon the emerging theory of abductive inference to establish a terminology for referring to abductive inference and its role in machine intelligence and learning.

(3) Use this terminology to translate the emerging theory of abductive inference into the terminology of state-of-the-art AI; use these translations to connect abductive inference research needs with current AI work that touches on abduction, e.g., nonmonotonic logic; and then discuss these connections within the AI community. (The point of such an exercise is to identify those aspects of current AI work which can contribute to the achievement of mechanized and autonomous abductive inference systems, and to identify a sequence of research steps that the AI community can take towards this goal.)

(4) Research proposals for specific machine intelligence projects should explain how the proposed project contributes to the ultimate goal of autonomous machine intelligence systems which learn by means of analytic, inductive, and abductive inferences. Enough is now known about the terms of this criterion to distinguish between projects which satisfy it and those which do not.

6.3 Natural Language and Other Man-Machine Communication

It is common sense that various specific communication goals are best served by different forms of exchange. This notion is borne out by the tendency in technical fields of human activity to spawn jargon which only slowly (if ever) filters into more widespread usage. In the general area of communication between man and machine, a few tasks are already well handled by available languages. For example, in the area of numerical computations the present formal languages, while not perfect, are highly serviceable.

When one considers the introduction of sophisticated computer systems into environments where it is necessary for them to communicate frequently, competently, and rapidly with people who are not specialists in computer programming, then the need for improvement in man-machine communication capability quickly becomes apparent.

A natural language capability in computers is required primarily in two kinds of circumstances — (1) where the nature of the information to be transferred warrants the flexibility and generality of a natural language, in distinction to a more specialized language, and (2) where, because of the number, nature, or condition of the humans involved, it is impractical to have the humans communicate in a formal language. There are additional considerations of convenience to the user, and attracting users who otherwise might be reluctant to use the available computer facilities.

For directing the global actions of robot devices of all kinds, as well as the interrogation of question-answering systems, the ability to use natural language considerably eases or eliminates the problem of training individuals to use these resources. In particular, the user population of an Earth-sensing information system can be significantly and economically extended through direct communication between users and the system in natural language. Unfortunately, at present the domain is essentially that of a research domain, with relatively few natural language interfaces operating in production environments.

Man-machine information exchanges can be segregated into iconic communication, such as pictures, and symbolic communication, such as formal computer languages and human natural language (see fig. 6.2). These differ significantly in the amount and kind of interpretation required to understand and to react to them. For instance, formal computer languages are largely designed to be understood by machines rather than people. For purposes of further discussion, man-machine communication is subcategorized as follows:

(1) Machine understanding of keyed (typed) natural language

(2) Machine participation in natural language dialogue

(3) Machine recognition/understanding of spoken language

(4) Machine generation of speech

(5) Visual and other communication (includes iconic forms).

6.3.1 Keyed Natural Language and Man-Machine Dialogue

In those instances in which the environment is highly restricted with respect to both the domain of discourse (semantics) and the form of statements which are appropriate (syntax), serviceable interfaces are just becoming possible with state-of-the-art techniques. Primary examples are the LUNAR (Woods et al., 1972) and SOPHIE (Brown and Burton, 1975) systems. However, any significant relaxation of semantic and syntactic constraints produces very difficult problems in AI. Intensive research is presently underway in this area. It seems that the semantic aspects of normal human use of language override a large part of the syntactical aspects. Computer languages traditionally have been almost entirely syntax-oriented, and so the considerable knowledge available concerning them has very little relevance in the natural language domain. Progress in flexible natural language interfaces is likely to be tied to progress in areas such as representation of knowledge and "common sense" reasoning. These lie at the heart of intelligent information processing — full natural language competency at the level of human performance requires a machine with intelligence and world knowledge comparable to that of humans. At this time there is little work in progress on the necessity or appropriateness of specialized hardware for natural language processing.

Accepting the close relationship between human-grade natural language proficiency and general intelligence level, and the improbability of near-term attainment of human-grade general intelligence in machines, it is appropriate to focus instead on achieving usable natural language interfaces at a lower level of machine performance. This leads to an examination of man-machine dialogues in which the goal of the man is to communicate a clear and immediate statement of information, or a request for information or action, to the machine, where the information or request is in a domain for which the machine has a competent model. In this sphere of activity the following component capabilities are thought to be highly desirable, and probably necessary, for efficient and effective communication: domain model, user model (general, idiosyncratic, contextual), dialogue model, explanatory capability, and reasonable default assumptions.
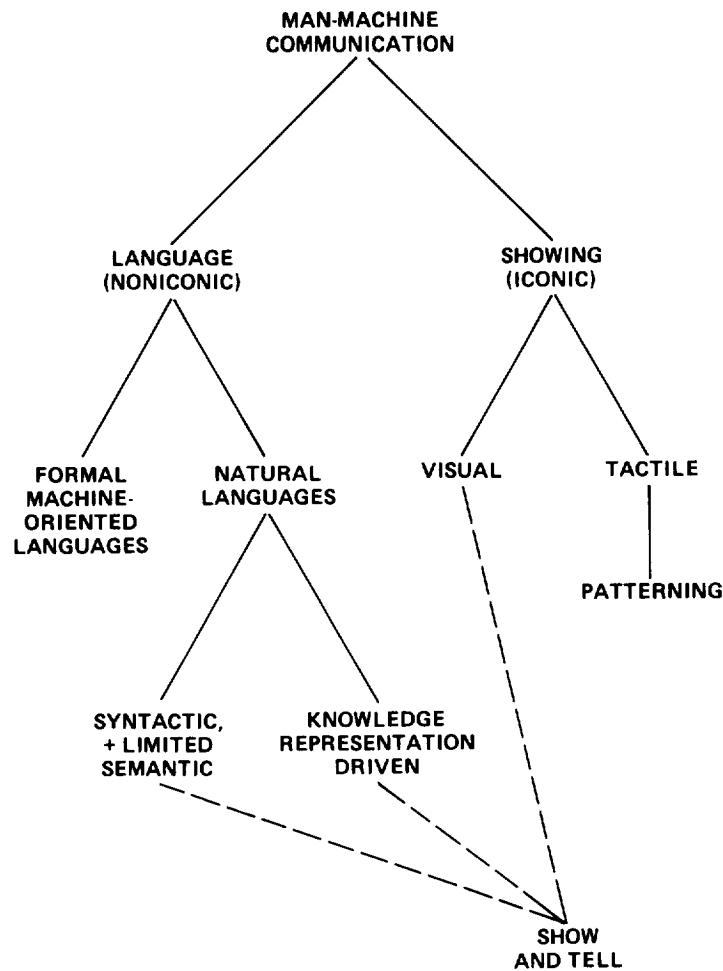
Figure 6.2.– Overview of man-machine communication.

Domain model. The machine must be able to act upon the information it receives, so it is assumed to have competency in some domain, called the object domain. To communicate about this object domain, the machine must have some additional knowledge called the domain model. When the communication environment is such that each linguistic transaction cannot have an immediate conclusion or effect in the object domain, then it is essential that the domain model be used to determine if a particular transaction makes sense. Otherwise the machine will not be able to make any inferences about the information it is being given, and the dialogue may become a monologue on the side of the human. The efficiency of the transfer of information will deteriorate, particularly for naive users. The machine will be accepting information which may prove inadequate later when it is handled in the object domain.

Individual user models. One of the reasons for natural language is to accommodate a wide range of humans in direct and efficient communication. This is best accomplished by taking into account at least some characteristics

which are either specific to particular individuals, or specific to classes of individuals. One example is in default values and assumptions. Different users will have different expectations concerning the values of implicit parameters of the conversation, and will have different underlying assumptions. In order not to burden each user with the necessity to make all of these explicit, it is necessary to make these assumptions and defaults a function of the type of user.

Dialogue model. The machine must have a working knowledge of what constitutes an acceptable dialogue. Such things as timing and absolute and relative explicitness are considerations pertinent to all users, and may vary from one person to another. In addition, the machine should avoid long series of questions posed to the human in order to clarify discussion, a particular problem with current expert systems such as MYCIN (Shortliffe, 1976). While part of the solution lies in proper default values and assumptions, there may still be need for the human to supply information in response to a perceived need by the

machine for this information. The best way to guide the human into supplying this information, while avoiding a tedious and long series of direct queries, is a largely unstudied area. Also, since the structure of the knowledge involved in the dialogue differs considerably between the human and machine, it will be necessary to map the initial internal "need to know" requests perceived by the machine into the general flow of the dialogue in a human-oriented way.

6.3.2 Machine Recognition and Understanding of Spoken Language

Recognition and understanding of fluent spoken language add further complexity to that of keyed language/phoneme ambiguity. In noise-free environments where restricted vocabularies are involved, it is possible to achieve relatively high recognition accuracy though at present not in real-time. In more realistic operating scenarios, oral fluency and recognition divorced from semantic understanding is not likely to succeed. The critical need is the coupling of a linguistic understanding system to the spoken natural language recognition process. Thus the progress in speech recognition will depend upon that in keyed natural language understanding.

Early applications have involved single word control directives for machinery that acts upon the physical world, using commands like "stop," "lower," "focus," etc. Some commercial equipment is available for simple sentences, but these require commands to be selected from a small predetermined set and necessitate machine training for each individual user.

6.3.3 Machine Generation of Speech

At the present time mechanical devices can generate artificial-sounding but easily understood (by humans) spoken output. Thus the physical aspects of speech generation are ready for applications, although some additional aesthetics-oriented technology work would be desirable. (The more important aspects of deciding what to say and how to phrase it were covered in the foregoing discussion of keyed natural language.)

6.3.4 Visual and Other Communication

Some motor-oriented transfer of information from humans to machines already has found limited application. Light pens and joysticks are rather common, and some detection of head-eye position has been employed for target acquisition. Graphics input/output is also an active research area, and three-dimensional graphical/pictorial interaction is likely to prove useful.

An interesting alternative approach in communicating information to robot systems is called "show and tell." In this method a human physically manipulates an iconic model of the real environment in which the robot is to act. The robot observes this action, perhaps receiving some simple coordinated information spoken by the human operator as he performs the model actions, then duplicates the actions in the real environment. The distinctions between show and tell and typical teleoperator modes of operation are:

- Show and tell does not assume real-time action of the robot with the human instruction.

- For show and tell, the robot has the time to analyze the overall plan, ask questions and generally form an optimal course of action by communicating with the human.

- The fidelity of the robot actions to the human example can vary in significant useful ways, allowing the robot to optimize the task in a manner alien to human thinking.

- The show and tell task can be constructed piecemeal, thus allowing a task to be described to the machine which requires many simultaneous and coordinated events, or which requires input from teams of human operators which is then chained together into a single more complex task description.

Show and tell permits a high degree of cooperative problem-solving and reasoning about actions between humans and machines. This novel technique probably has an important functional role to play somewhere between autonomous robots and pure teleoperation.

6.3.5 Recommendations

The team recommendations to NASA regarding directed research and development in the field of natural languages and other man-machine dialogue are as follows.

Natural language and man-machine dialogue. Theoretical work in keyed and spoken natural language for managing restricted domain databases will proceed with NASA involvement. The impact of such systems is widely recognized, and much research is in progress. In applications, DOD is already involved in funding research whose results will probably be directly applicable to NASA database interactions in the immediate future. It is recommended that NASA now make plans to initiate implementation of systems using keyed natural language for internal use within NASA. Such implementation not only will provide useful production tools for NASA, but also will generate the in-house experience necessary to provide these techniques to outside users of space-acquired data as in the IESIS mission.

More sophisticated uses of natural language, such as in directing almost autonomous robots in tasks like space

354

construction or exploration, should be studied by NASA, but immediate application of a pure natural language communication channel does not seem possible at this time or in the near future. The first uses of fluent natural language in controlling robots will probably best be done in a context such as "show and tell."

Spoken language. The development of fluent spoken language recognition is expected to evolve in step with the ability of machines to understand and reason about the object domain. Thus, a NASA orientation toward funding research in this area would be misdirected. There is no obvious pressing need at this time for the Agency to intervene in the development of isolated word recognition control of robots, as this area will develop very rapidly on its own.

Speech generation. Serviceable speech generation is technologically current, for the physical generation, and NASA need not take any particular steps in this area until specific implementation demands it. The more important area of machine decision of what to say is a much more difficult and undeveloped research area, and is essentially the same problem as in keyed input-visual output dialogue systems.

Visual and other communication. The areas of motor and graphic interaction are ready for current implementation. NASA should consider these as tools appropriate both for its own internal use and, as with the keyed natural language, for outside users of NASA-collected data. Show and tell communication would be extremely useful in zero-g robot-assisted construction, and may have application in planetary exploration and space or lunar industrial processes – current research efforts are minimal. Many of the specific capabilities of potential interest to NASA will not be developed if the space agency does not take a direct, active role.

A very rudimentary form of show and tell, called "patterning," should be implemented as soon as possible for all NASA spacecraft with manipulator or other movable components under computer control. In patterning, a prototype or other model of the actual spacecraft is physically articulated in the way the actual spacecraft should behave. The model is connected to a computer through appropriate proprioceptors, and the computer writes a program which can be uploaded to the spacecraft to direct its actions. It should also be required that the model be able to execute the program in order to verify its correctness. Such a capability would greatly extend the flexibility of control of both complex devices in space and exploration craft on planets, and yet are relatively easily implementable with current techniques.

6.4 Space Manufacturing

To achieve the goal of nonterrestrial utilization of materials and factory self-replication and growth, space manufacturing must progress from terrestrial simulation to low Earth orbit (LEO) experimentation with space production techniques, and ultimately to processing lunar materials and other nonterrestrial resources into feedstock for more basic product development. The central focus of this assessment is upon the technologies necessary to acquire a major space manufacturing capability starting with an automated Earth orbiting industrial experimental station established either as an independent satellite or in conjunction with a manned platform such as a manned orbiting facility or "space station."

6.4.1 Earth Orbiting Manufacturing Experiment Station

There are four major components of any production system: (1) extraction and purification of raw materials, (2) forming of product components, (3) product component assembly, and (4) system control. The Earth-orbiting station will conduct experiments to determine the relative merits of alternative methods of implementing these elements in a space manufacturing facility.

Product formation involves two general operations – primary shaping to achieve the approximate shape and size of the component and finishing to meet all surface and dimensional requirements. The most promising primary shaping technologies for space manufacturing are casting and powder-processing techniques. When properly controlled, both methods produce parts ready for use without further processing. Casting techniques appear more versatile in terms of the range of materials (metals, ceramic, metal-ceramic) that can be shaped, but powder processes may outperform casting for metallic components. A determination of the relative utility of these two processes should be one of the primary goals of the space manufacturing experiment station.

The casting process is a fairly labor-intensive activity on Earth and has not been highly automated, with the exception of Strand and other continuous casters. Automated casting facilities do not generally produce a variety of parts configurations; instead, they usually make just a single shape (usually a bloom or billet) which later requires a great deal of expensive and time-consuming processing before it is usable as a machine component.

Many of the finishing operations can be eliminated if the material is cast into (approximately) its final configuration using a specialized mold. The production of these molds has been automated in two instances. In investment casting, the dipping of the wax forms into ceramic slurry has been

accomplished using industrial robots, although actual pattern formation remains largely manual. In permanent die casting, the Nike sports shoe subsidiary in Massachusetts produces tapes for its N/C electrical discharge machining apparatus which drives the tool to form the dies automatically from drawings of the shoe's sole constructed using the plant's CAD/CAM system.

On the whole, however, the formation of patterns and disposable molds (especially green sand casting molds) has remained manual; only equipment for lifting and turning the flasks has come into widespread use. Robots have been employed to unload hot parts from die casting machines as well as place the (hot) castings into trimming dies. Almost all automation of casting has been in high-volume applications where one standard shape is produced ten thousand times per year or more. High-volume production is not likely to be the general mode of space manufacturing, which will probably call for small lot, intermittent production. Methods of performing automatic casting, especially using disposable molds, and doing this efficiently in low- or zero-gravity conditions are required. Elimination of molds using containerless forming techniques should also be investigated and, if successful, will significantly reduce the high capital costs of forming molds and dies. The problems of heat removal and control of the rate of cooling to control grain size in the castings requires both sensor development to sense the internal temperatures and new heat dissipation technologies.

Powder processing has been somewhat automated on Earth, but has not been used extensively due to the tremendous costs associated with purifying and maintaining a non-oxidizing environment for manufacturing. This environment is available in space and on the lunar surface. But, as in the case of casting, powder processing uses dies to form parts. Again, the study of containerless forming techniques may be fruitful, with powder processing alleviating some of the heat dissipation problem, since sintering temperatures are lower than those required for casting. The applicability of powder formation via liquefaction and spraying should be assessed. Grinding and milling must also be examined, since the cold welding phenomenon between similar pure metals may be turned to advantage if it can be used to facilitate coalescence of the metals without sintering or melting. Intensive study of this effect is best performed in space, as pure powders are extremely difficult to prepare and maintain on Earth. Cold welding also has important implications for machining and lubrication.

Machining, or chip formation processes, are the usual finishing operations. These have been extensively automated, but significant problems with heat dissipation and cold welding may be encountered in space if the tools are run in a vacuum. The primary cause of tool wear is the temperature generated at the tool/chip interface. Removal of this heat through the use of cutting fluids will be difficult because all terrestrially used fluids are either petroleum

or water-based — two commodities expensive in space and difficult to control in a zero-g environment. Cold welding will decrease chip forming in two ways — first, by the formation of built-up edge on the tool face (although temperature and pressure may still be the determinants of this effect), and second, by the reattachment of the pure metal chips to the cut or uncut surface or machine table by vacuum welding. Use of lasers to finish may eliminate many of these problems and thus may be of tremendous utility, especially if casting or powder techniques can be expected to produce high-tolerance parts. The use of ultra-high speed machining in which most of the heat of cutting is carried away by the molten chip could also be a partial solution, and also the use of ceramic tool bits and cast basalt tables. (See also appendix 5F.)

Assembly requires robotic/teleoperator vision and end-effectors which are smart, self-preserving, and dexterous. Accuracy of placement to 0.001 in. and repeatability to 0.0005 in. is desirable for electronics assembly. Fastening technologies, including nonvolatile adhesives, cold welding, mechanical fasteners, and welding all require special adaptation to the space environment.

Control of a large-scale space manufacturing system demands the use of a distributed, hierarchical, machine-intelligent information system. Material handling tasks require automated, mobile robots/teleoperators. In support of these activities, vision and high-capacity arms, multi-arm coordination, and dexterous end-effectors must be developed. For inventory control, an automated storage and retrieval system well suited to the space environment is needed. The ability to gauge and measure products (quality control) benefits from automated inspection, but a general-purpose machine-intelligent high-resolution vision module is necessary for quality control of complex products.

6.4.2 Materials Processing and Utilization

While it is expected that the orbiting space manufacturing experiment station initially will be supplied with differentiated raw feedstock for further processing, some interesting experiments in systems operations and materials extraction are possible and should be vigorously pursued. One such experiment could be a project to build one reasonably complex machine tool using a minimum of human intervention and equipment. Two logical candidates emerge. The first is a milling, grinding, or melting device that could be used to reduce Shuttle external tanks to feedstock for further parts building or experiments. This project would allow experimentation in material separation and processing using a well-defined and limited input source which can be obtained at relatively low cost whenever the Space Shuttle carries a volume-limited rather than a weight-restricted load. Such a large-scale experiment could be used as an "extra-laboratory" verification of extraction,

manipulation and control mechanizations, as well as providing relatively easy access to pure metal powders for research. A second candidate project is the fabrication and assembly of a beam-builder for use in large structure construction experiments. These two machine tool projects could then be combined to study materials handling and storage problems by having the first project provide feedstock for the second. Additional experimentation on producing feedstock from lunar materials would be a logical outgrowth of this development.

While the space manufacturing experiment station is largely viewed as an experiment station for capital equipment production and as a stepping stone to the establishment of a lunar manufacturing facility, it should be noted that the station can also be used for biological research and the preparation of products such as drugs and medicines for terrestrial consumption. For example, many pharmaceutical components require a zero-g environment for their separation. Additional products for terrestrial consumption would be perfect spheres or flat surfaces made by joining bubbles.

The technology required for permanent facilities to process nonterrestrial materials on the lunar surface or elsewhere lies far beyond currently proposed space materials processing capabilities. Numerous workers have suggested processes such as electrolysis, hydrogen fluoride leaching and carbochlorination (see section 4.2.2), which are adequate for short-term usage but cannot reasonably be expected to meet long-term growth requirements. Processes must be developed which yield a far broader range of elements and materials, including fluorine, phosphates, silica, etc. Volatiles such as water and argon, and desirable rock types such as alkalic basalts and hydrothermally altered basalts, could be acquired as a result of lunar-surface exploration. High-grade metals can probably be retrieved from asteroids in the more distant future.

Sophisticated and highly automated chemical, electrical, and crystallization processing techniques must be developed in order to supply the wide variety of required feedstock and chemicals. Some possible solutions may be generated by studying controlled fractionation and chemical doping of molten lunar materials in order to achieve crystallization of desired phases. Zone refining and zone melting techniques may also be fruitful areas for investigation. New oxygen-based chemical processing methods should also be examined.

6.4.3 Technology Requirements

The control of individual machine tools has continued to advance in feedback and feedforward control modes. The control of a diverse, highly integrated industrial complex requires advances in computer systems. High-speed data access in linked hierarchical computer networks will be needed. These computers will require coordination in real time. For example, the material handling computers must

relay messages to the material handling devices telling them which machines need to be emptied or loaded and the material handling devices must know where to place the removed product. Advances in autonomous planning and scheduling in a dynamic environment are required, using new scheduling algorithms and shop floor control techniques. Large database requirements will soon become apparent. Repair robots must have the capability to hypothesize probable causes and sources of malfunction.

The establishment of space or lunar manufacturing facilities require the development of the following technologies:

- Basic research on materials processing in the space environment

- Improvement in primary shaping technologies of casting and powder processing for metals and nonmetals with emphasis on the economic elimination of manual mold production, possibly by the use of containerless forming

- Improvement in heat dissipation abilities in relation to the tool/chip interface in space, and control of cooling rates in castings

- Comprehension of cold-welding as a limiting factor for metal curing and as a joining technique

- Improvement of robot dexterity and sensors (especially vision)

- General and special purpose teleoperator/robot systems for materials handling, inventory control, assembly, inspection, and repair

- Improvement in computer control of large, integrated, dynamic hierarchical systems using sophisticated sensory feedback

- Study and improvement of lasers and electron-beam machining devices

- Embodiment of managerial skills in an autonomous, adaptive-control expert system

6.5 Teleoperators and Robot Systems

A teleoperator is a device that allows action or observation at a distant site by a human operator. Teleoperators represent an interim position between fully manned and autonomous robot operation. Teleoperators have motor functions (commanded by a human) with many possible capabilities, and have sensors (possibly multiple, special-purpose) to supply information. The human being controls and supervises operations through a mechanical or computer interface. As technology advances and new requirements dictate, more and more of the command and control functions will reside in the computer with the man assuming an increasingly supervisory role; as artificial intelligence methods are developed and are applied, the computer eventually may perform "mental" functions of greater

complexity, making the system more autonomous. The following discussion concerns teleoperators and their functions, applications to NASA programs, necessary supporting technologies, and the evolutionary path of robotics.

Teleoperators have been developed to expand man's physical capabilities across great distances and in hostile or inaccessible environments. Typical applications include (1) safe, efficient handling of nuclear or toxic materials, (2) undersea mining and exploration, (3) medical and surgical techniques, and (4) fabrication, assembly, and maintenance on Earth and in space. An artificial limb is considered a teleoperator because it restores lost dexterity to an amputee.

Teleoperators are not new. In 1954 Argonne National Laboratory developed a master/slave hand system with force feedback via cables and pulleys. In 1958 William Bradley (1980) operated an area-of-interest television camera system mounted on a truck to provide a display to the "driver" located 15 km away. In the 1960s General Electric engineers designed "Hardiman," an exoskeletal teleoperator with 15 degrees of freedom and the capability of manipulating 700-kg loads with ease (Corliss and Johnsen, 1968). Research is progressing once again in manipulators, sensors, and master/slave systems. Further technology advances will be made as NASA develops teleoperators for space operations.

6.5.1 Teleoperator Applications

Advanced teleoperators for future space missions present new challenges in the development of spaceborne man-machine systems (Bejczy, 1979; Bradley, 1967; Corliss and Johnsen, 1968). Teleoperators are robotic devices having video or other sensors, manipulator appendages, and some mobility capability, all remotely controlled via a telecommunications channel by a human operator. The man can exercise direct in-the-loop control using a joystick or other analog device, or can choose more indirect means of command such as an AI system in which he shares and trades control with a computer (NASA Advisory Council, 1978). Heer (1979) estimates that flight demonstrations of automated Shuttle manipulators can begin as early as 1982, for automated construction devices in 1986, and for a free-flying automated teleoperator by 1987.

A teleoperator will be on the first operational Space Shuttle flight. The Shuttle has a six-degree-of-freedom general-purpose Remote Manipulator System (RMS) with a 15-m reach (Meade and Nedwich, 1978; Raibert, 1979). The RMS lifts heavy objects in and out of the payload bay and assists in orbital assembly and maintenance (Meade and Nedwich, 1978; Raibert, 1979). An astronaut controls the rate of movement of the RMS using two three-axis hand controllers (Lippay, 1977). One proposed follow-on is installation of a work platform so that the RMS could be used as a "cherry picker," carrying astronauts to nearby work sites.

One RMS will be mounted on the port longeron with provisions for a second RMS mounted on the starboard side. Conceptually the RMS arm is much like a human arm, with yaw and pitch at the shoulder joint, pitch at the elbow, and yaw, pitch and roll at the wrist (Lippay, 1977). The upper arm is 6.37 m long and the lower arm is 7.06 m long, providing a 15-m reach. The RMS can move a 14,000-kg payload at 6 cm/sec with the arm fully extended, or up to 60 cm/sec with no load (Space Shuttle, 1976, 1977, 1978).

Two other distinct classes of teleoperation will be required for complex, large-scale space operations typified by the space manufacturing facility described in chapter 4. The first is a free-flying system which combines the technology of the Manned Maneuvering Unit with the safety and versatility of remote manipulation. The free-flying teleoperator could be used for satellite servicing and for stockpiling and handling materials (Schappell et al., 1979). Both of these operations require autonomous rendezvous, stationkeeping, and attachment or docking capabilities. Satellite servicing requires the design of modular, easily serviceable systems and concurrent development of teleoperator systems.

The Teleoperator Maneuvering System (TMS) is an unmanned free-flying spacegoing system designed to fit in the Shuttle Orbiter, with the capability to boost satellites into higher orbits, service and retrieve spacecraft, support the construction assembly and servicing of large space platforms, capture space debris, and perform numerous other tasks in orbit. TMS has the potential, with developing robotics technology, to greatly extend and enhance man's capabilities in space. As presently defined by NASA, TMS is propelled with hydrazine or cold gas thrusters, controlled by operators at ground stations or in the Orbiter's aft flight deck, and can be placed under automated control using its onboard computational capabilities. TMS eventually will be equipped with antennas, manipulators, video equipment, dexterous servicing mechanisms, a solar power array, and other equipment as needed to position spacecraft, rendezvous with and service satellites, position large platform sections, and act as a "smart" free-flying subsatellite for performing specialized missions. It can perform all known LEO payload retrieval missions within 1 km of the Shuttle, and retrieval at distances of 800–1600 km from the Orbiter could be demonstrated by the mid-1980s (OAST, 1980).

Manufacturing processes and hazardous materials handling may utilize mobile or "walking devices," the second distinct class of teleoperators. The teleoperator would autonomously move to the desired internal or external site and perform either preprogrammed or remotely controlled operations. For manufacturing and repair, such a system could transport an astronaut to the site and the manipulator could be controlled locally for view/clamp/tool operations or as a workbench. Of course, the size and level of teleoperator mobility (free-flying or walking) is dictated by mission needs.

Probably the two classes will be combined into one device in actual practice. Such a combination could be characterized as a remote free-flying teleoperator equipped with a highly specialized manipulator of the general-purpose or spacecraft-services system type. An extension can be envisioned as a teleoperator vehicle combined with single or multiple manipulator arms used to align and attach beams in direct support of space construction activities. Depending on the complexity of the task at hand it may be necessary for humans to be directly in charge in a master-slave relationship and be housed in a life support module on the free-flyer. The next logical development step delegates this function to a manlike robot, thus freeing the system to work autonomously at extended operational ranges without the cumbersome remote or local presence of man.

In the reference Space Manufacturing Facility developed by Miller and Smith (1979), the large number of similar components in the solar-cell factory and the X-ray environment precludes direct human labor. This suggests automated maintenance and repair, so the solar-cell factory was designed for tending by automated and remote devices. A free-flying hybrid teleoperator (FHT) can do on-site repairs at the solar-cell factory. The FHT can be operated either fully automated (tied into an AI-capable computer system or using preprogrammed routines), automatically with human override, or fully remote-controlled by a human operator (teleoperation).

Free-flying teleoperators or robot servicing units will have the capability to autonomously rendezvous, close, and attach to a satellite, first in LEO near the main station and later in GEO (Schappell et al., 1979). In some cases satellite retrieval, rather than servicing, will be desired. This would be a precursor to automated asteroid retrieval missions, requiring completely autonomous systems for navigation, guidance, sensing and analysis, attachment, and mining (Shin and Yerazunis, 1978).

On-board and free-flying teleoperators will be required throughout the postulated mission plan. They will extend man's senses and dexterity to remote locations while the human supervises and controls from a safe, comfortable environment. Teleoperators are a logical step in the evolution to fully automated (robot) systems needed for efficient extraterrestrial exploration and utilization. Previous sections have already discussed the role of man, the role and configuration of such teleoperators, and the role and development required for completely automated, possibly self-replicating, systems.

6.5.2 Teleoperation Sensing Technology

The uniqueness and utility of teleoperators lies not in their mode of locomotion, but rather in the "telepresence" they provide — the ability of the man to directly sense and remotely affect the environment (Minsky, 1980). Sensor and manipulator technology is advancing apace, largely through rapid growth in the fields of industrial robotics and computer science.

Approximately 40% of human sensory input is in the form of vision, so it is appropriate that most work in physical perception relates to visual information processing and remote scene interpretation. Algorithms and specialized sensors developed for satellite on-board pattern recognition and scene analysis can enable the teleoperator to perform many of these functions. Teleoperation has several unique characteristics such as viewing and working in three dimensions under variable conditions of scene illumination, and options of wide or restricted fields of view. Three-dimensional information can be obtained from stereo displays (Chin, 1976; Duda and Hart, 1978; IEEE, 1979), lasers (Shin and Yerazunis, 1978), planar light beams (Baum, 1979), radar and proximity sensors (Schappel et al., 1979), or it may be recovered from two-dimensional pictures (Tenenbaum, 1979).

Besides its use in autonomous tasks, a computer "world model" can be utilized in two ways. First, it can provide the man a computer-generated display from any point in the "world." Theoretically, from an overall view of the entire scene (including the teleoperator itself) the camera eye could zoom down inside a crevice or behind an object. Using data from a scanning laser ranging system, the system described by Shin and Yerazunis (1978) could construct a perspective model of nearby terrain and superimpose the route through the terrain determined by an optimal path-selection algorithm.

Second, using camera location as a reference point and overlaying the "world model" over the camera picture would permit correlation of the world model with the real world, thus enabling the operator to immediately detect anomalies or inaccuracies in the knowledge base. This "knowledge overlay" would allow corrections for sensor errors and keep autonomous manipulator operations properly referenced. Without such a knowledge overlay the man is severely handicapped in acting as supervisor of largely autonomous operations.

Besides vision, a teleoperator should give the human a "feel" for the task. Minsky (1980) notes that no present system has a true sense of feel, and insists that "we must set high objectives for the senses of touch, texture, vibration, and all the other information that informs our own hands." In addition to communicating via sight and touch, an audio interface between man and computer also is feasible (see section 6.5.3). Voice input/output systems are commercially available and in use. Research continues, though, in artificial intelligence and computer science on natural language understanding, faster algorithms, and connected speech processing. However, it should be noted that teleoperators with simple bilateral force reflection can achieve most immediate goals in space. These were demonstrated by Ray Goertz as early as 1955, and can be used now.

6.5.3 Teleoperation Manipulator Technology

Much of a teleoperator's capability is sensory; much is associated with manipulation. Although configurational details require further definition of task requirements, overall general-purpose space teleoperator characteristics can be partly inferred. A teleoperator arm must have enough freedom so that the manipulation and arm locomotion systems can position the hand or end-effector at any desired position in the work environment. There must also be a locus of points which all of the teleoperator's hands can reach simultaneously. If such a region does not exist, manipulator cooperation is precluded — cooperation and coordination of multiple manipulator arms and hands give teleoperators (and humans) tremendous potential versatility.

How many manipulator arms might the general-purpose teleoperator have? Despite man's two arms, the teleoperator will probably need three. Most mechanical operations require just two hands — one to grasp the material and the other to perform some task. A third hand would be useful in holding two objects to be joined, or in aiming a television camera (or other appropriate sensor). In many two-handed operations on Earth the human worker moves his head "to get a better look" — the third teleoperator arm would move the man's remote eyes for that purpose. Indeed, the third arm can be used to couple the TV motion to the man's head motion. Bradley (1980) notes that this gives a strong feeling of telepresence. Finally, three fingers probably are sufficient for duplicating most of the functions of the human hand — this is the minimum number necessary for a truly stable and controllable grasp of small objects.

6.5.4 Robot Systems

Teleoperators will always be vital to many operations in space because they extend man's senses and motor functions to remote locations. But extraterrestrial exploration and utilization and other advanced systems will require remote autonomous systems — systems with on-board intelligence. These robot systems will evolve along with current AI efforts at representing knowledge functions in a computer. The integration of AI technology with teleoperator/robot systems is a major development task in its own right and should be timed to support space programs that require this capability.

Aspects of artificial intelligence which must be addressed in regard to robot systems include memory organization, knowledge retrieval, search, deduction, induction and hypothesis formation, learning, planning, perception and recognition (Lighthill, 1972; Nilsson, 1974; Sagan, 1980; Winston, 1978). Teleoperation and robotics technology requirements are: time lag compensation methods, sensory scaling, adaptive control methods, touch sensing, hands, hydraulics, actuators that are many times lighter than the

masses that they lift, onboard power for autonomous operation (this is a major problem), parallel computers, clamp and hold servoing of arms (extra hands are needed to hold parts while soldering and connecting), homeostasis, survival instincts, world models, laser data links, and laser sensors. Computer science, cybernetics, control theory and industrial process control are all relevant fields in this research. Interactive systems are being developed whereby the computer works, not autonomously, but as a partner or intelligent assistant. Kraiss (1980) discusses the design of systems resulting from cooperation of human and robot systems in four specific areas — computers capable of learning and adapting, computer support in preparation and evaluation of information, computer support in decisionmaking, and computer assistance in problem-solving.

6.5.5 Telefactor Technology Development Recommendations

The advantages of the availability of telefactor systems for development of subsequent fully automatic and replicating systems have already been described in this report. However, it is worth noting that: (1) all of the technical information and components to build a telefactor system were available, and the basic subsystems (e.g., master-slave manipulators and head-aimed television systems built and demonstrated) before 1965, and (2) to date, no one has built a complete system (Bradley, 1967).

Construction of a standard telefactor system is long overdue. NASA should include this important step in an early phase of its automation program. Some of its applications to the NASA program are the following:

1. A telefactor system can be used to oversee and operate a materials processing activity to establish requirements for full automation of such activity and also for manned intervention.

2. A telefactor system can provide a built-in maintenance and repair facility in a complex spacecraft.

3. A telefactor system could perform satellite inspection, modification, or other EVA operations from the Shuttle, even with uncooperative objects.

4. All of the actions and observations of a telefactor system can be taped for later playback, permitting retrospective task analysis.

5. Demonstration of the frequently proclaimed versatility and effectiveness of telefactor systems is overdue and much needed.

6. A standard telefactor system can be used as a comparison-piece in the field of robotics. Differences in task performance and in characteristic deficiencies among telefactors versus robots would be of great interest.

7. Since computers can be readily inserted into a standard telefactor system, these could become powerful tools in the development of fully automatic or supervisory control systems.

8. A standard telefactor system would be a convenient starting point for development of rovers for lunar and planetary exploration and prospecting.

9. A standard telefactor system would incidentally be a useful starting point for development of terrestrial remote control or remotely piloted vehicle equipment for use in hostile environments.

To achieve construction of a prototype standard telefactor system with minimum cost, it would appear appropriate to utilize some of the personnel already familiar with this and who have had practical experience in the construction and operation of the major subsystems. A conventional aerospace contractor even if well provided with funding and facilities is likely to misunderstand some of the problem areas discovered and resolved during 1956-1966, thus requiring costly reworking and rediscovery of old techniques. NASA should find means of implementing such an effort with leadership at one of its centers and with interested participation by NASA Headquarters staff. After attainment of a satisfactory prototype of a standard telefactor system, several should be constructed and made available where needed most in the agency's automation program.

6.6 Computer Science and Technology

NASA's role, both now and in the future, is fundamentally one of information acquisition, processing, analysis, and dissemination. This requires a strong institutional expertise in computer science and technology (CS&T). Previous study efforts and reports have made recommendations to integrate current technology more fully into existing NASA programs and to develop NASA excellence in selected relevant fields of computer science. Recent studies have explored the research and development requirements of NASA field centers, and have identified particular R&D goals and objectives relevant to CS&T (ASC, 1980; EER, 1980; Sagan, 1980). In this section, the recommendations are considered from the perspective of the CS&T study team, together with the implications for CS&T of the various missions defined earlier in the report.

Of particular concern in the present technology assessment is the evolving CS&T program required within the space Agency to support a major involvement of automation and machine intelligence capabilities in future NASA missions. The agency presently is not organized to support such a vigorous program in CS&T. Most apparent is the lack of a discipline office at the Headquarters level which supports research and development in computer science and which serves as an Agency advocate for the incorporation of state-of-the-art capabilities into NASA programs.

NASA technical requirements with relevance to CS&T are presented and correlated with specific CS&T disciplines in this section. A general upgrading of computing facilities is recommended. Building an organization to maintain a state-of-the-art capability in the computing and information sciences is perhaps the greatest challenge for the future. The study group is hardly qualified to offer specific organizational recommendations to NASA, but encourages the agency to consider an organizational response and suggests some ideas which may be helpful. Finally, maintenance of a solid computing science institutional capability depends on a vigorous and continuing program of intellectual exchange with peer organizations in academia, industry, and government. A few suggestions are presented as to the possible components of such a program.

6.6.1 NASA Technology Requirements

This report, together with the report of the NASA Study Group on Machine Intelligence and Robotics (Sagan, 1980), has explored the application of advanced automation within NASA. In addition, there are general computer science capabilities required to develop and implement the types of missions described in the present document. These include robotics, smart sensors, mission operations, computer systems, software, data management, database systems, management services, human-machine systems, engineering, and system engineering.

Robotics. The principal requirements associated with robotics which call upon the disciplines of CS&T include visual perception, manipulator control, and autonomous control. This latter category includes problem-solving and plan generation - the ability of a robot device to plan and to pursue its own macroscopic course of action. NASA requirements also argue for a robotic capability to perform intelligent data gathering and in some instances to provide a telepresence capability for a remote human operator.

Smart sensors. Current programs such as NEEDS (see chapter 2) address requirements for smart sensing devices which selectively acquire data and analyze it for information value prior to consuming communications and storage capacity. These requirements include visual perception, image processing, pattern recognition, scene analysis, and information extraction. In addition, the notion of model-based sensing shows promise for intelligent data acquisition. To conserve communications bandwidth, user-oriented data compression techniques are required which can result in a several orders-of-magnitude reduction in the amount of data transmitted.

Mission operations. In the area of mission operations, a rather general symbolic modeling and representation capability is required to do planning, scheduling, sequencing, and monitoring, as well as fault modeling and diagnosis. This draws on problem-solving techniques within artificial

intelligence and can benefit greatly from a hypothesis formation capability. As mission operations are presently conducted, machine intelligence can benefit the coordination of manpower, as well as enhance the mission software development and integration process. As mission operations are envisioned in the future, involving autonomous spacecraft operations and automatic mission control, a strong dependence on CS&T in general and machine intelligence in particular is unavoidable.

Computer systems. In ground-based systems, and especially in spaceborne applications, NASA has a fundamental dependence on computer systems. Requirements include LSI and VLSI circuit design, fabrication and test techniques as well as fault-tolerance, error detection and recovery, component reliability, and space qualification. Beyond the component level, very significant primary and secondary storage requirements emerge. System-level issues become dominant, such as computer architecture (e.g., parallel processors) and system architecture (e.g., computer networks). Many of NASA's systems have severe real-time constraints, and techniques for adequate system control demand attention.

Software. Much of NASA's technology resources are spent on software, yet relatively modest attempts have been made to improve the process of software development, management, and maintenance. Given the exciting prospects for computer-based advanced automation in future missions, a program for more efficient, effective, and timely software development, management, and maintanance is mandatory. Principal software requirements are in the areas of programming languages, the software development environment, software validation, algorithm design, fault tolerance, and error recovery. Automatic programming should also be considered as a vehicle for improving the quality of software and the process of developing it.

Data management. Data management requirements comprise a very large part of the CS&T-related requirements within NASA, and include most of the interfaces to the user community. The public perception of NASA's systems will be derived largely from their ability to use them and to derive benefit from them. Both the NEEDS and ADS projects have realized this, and have diligently considered the end-user interface requirements. Data management requirements include data compression, staging, integration, and dissemination, as well as the implied requirements of data autonomy. Scheduling, performance monitoring, and system control also imply data management requirements, as does sensor management. A fundamental element of a user-oriented system is an extensive directory service as well as a capability to model the user, to know the context of his requests and his level of sophistication. On-line tutorial capabilities are

appropriate for a diverse user community, and provide valuable input to the development of a user model. Knowledge-base systems and constructs such as semantic networks can contribute greatly to NASA's data management capability.

Database systems. NASA's current database requirements are not considered to be extraordinary from the CS&T perspective, although future systems supporting a geographically dispersed, technically diverse user community attempting to analyze or correlate sets of data spanning several distinct databases will require a sophisticated capability currently beyond the state-of-the-art. The requirements in this area include "traditional" database systems as well as relational database systems. The capability to satisfy queries which require access to several geographically separate databases is considered fundamental, as is a complete archiving capability.

Management services. NASA has, to a large extent, avoided the application of contemporary CS&T (let alone machine intelligence) to the management of the agency itself and its own programs. Current commercial offerings in management information and word processing systems can substantially enhance the efficiency and effectiveness of NASA management, both at Headquarters and at the field Centers. State-of-the-art capabilities in on-line records management, calendar coordination, and "bulletin boards" can likewise have a significant positive impact. The automated office is a concept evolving from this work which could revolutionize NASA's management techniques. Some obvious requirements in this area are manpower coordination, document preparation, and forms processing (e.g., travel orders and procurement requests). Presently unexplored is the potential application of contemporary machine intelligence techniques such as problem-solving, reasoning, and hypothesis formulation to the management of projects and the exploration of policy alternatives.

Human/machine systems. NASA has extensive requirements relating to human/machine interactions and currently has several efforts exploring the application of machine intelligence to these problems, primarily in the areas of hand-eye coordination and natural language processing. Requirements are primarily in the areas of human/ machine control processes and the interface between a human and an "intelligent" computer system. Coordinated work between the computing sciences and cognitive psychology may be required to make substantial progress in this field.

Engineering. NASA is currently applying state-of-the-art technology in the engineering disciplines, particularly in computer-aided design, manufacturing and testing. The requirements of future missions, including mining and

manufacturing in nonterrestrial environments, mandate a continuing vigorous program in this area, embracing robotics technology.

System engineering. There are many component technologies which must come together to build a system. CS&T can aid in the process of engineering systems solutions, rather than component solutions, to systems problems. Formally managing the definition of requirements for a system is one example. Other contributions of CS&T include formalized design methodologies, techniques for performance monitoring and evaluation, and quasi-rigorous approaches to system architecture and control. Requirements in each of these areas pervade NASA programming.

6.6.2 Relevant CS&T Disciplines

For purposes of the present discussion the scope of CS&T is considered to be that classified by the Association for Computing Machinery (Computing Reviews, 1976) in their recently published "Categories of the Computing Sciences." This basic taxonomy was reviewed by the CS&T study team. Those components which appeared to relate most strongly to NASA's anticipated future requirements were scrutinized in more detail. The results of this analysis are summarized briefly below, with the ACM classification number included parenthetically for completeness.

Applications (3.). "Applications" focuses on the uses of computers and the relationships between human cognitive and perceptual processes and computers. NASA-relevant subcategories include:

(3.1) Natural sciences (astronomy, space, earth sciences)

(3.2) Engineering (aeronautical, electronic, mechanical)

(3.4) Humanities (language translation, linguistics)

(3.5) Management (policy analysis, manufacturing, distribution)

(3.6) Artificial intelligence (induction, pattern recognition, problem-solving)

(3.7) Information retrieval (content analysis, file maintenance, searching)

(3.8) Real-time systems (process control, telemetry, spacecraft simulation)

Software (4.). This category includes "the procedures, instructions, techniques, and the data required to apply a computer to a given task." Relevant subcategories include:

(4.2) Programming languages (procedure, problem-oriented)

(4.3) Supervisory systems (multiprogramming, database systems)

(4.4) Utility programs (debugging, program maintenance)

(4.6) Software evaluation, test, and measurements (software modeling, algorithm performance monitoring)

Mathematics of computation (5.). The category of mathematics of computation consists of "the intersection of mathematics and computer science, the category embraces subcategories that cover the mathematical treatment of numbers, mathematical metatheory, symbolic algebraic computation, the study of computational structures (algorithms, data structures) as mathematical objects, and mathematical methods that lend themselves to computer-aided solutions." This entire category was considered relevant to NASA. The subcategories are:

(5.1) Numerical analysis (error analysis, numerical integration)

(5.2) Metatheory (logic, automata, formal languages, analysis of programs)

(5.3) Combinatorial and discrete mathematics (sorting, graph theory)

(5.4) Mathematical programming (linear and nonlinear programming, dynamic programming)

(5.5) Mathematical statistics and probability (regression and correlation analysis, stochastic systems)

(5.6) Information theory (decision feedback, entropy)

(5.7) Symbolic algebraic computation (symbolic differentiators, symbolic interpretors)

Hardware (6.). The hardware category includes all of the physical components of digital computers. The relevant subcategories include:

(6.1) Logical design, switching theory (functional design, switching networks, Boolean algebras)

(6.2) Computer systems (packet switching networks, time-shared hardware, parallel processors)

(6.3) Components and circuits (LSI/VLSI, control and storage units)

Functions (8.). This category deals with major computer functions and techniques. NASA-relevant subcategories include:

(8.1) Simulation and modeling (applications, techniques, theory)

(8.2) Graphics (display processors, image processing, plotting)

(8.3) Operations research/decision tables (PERT, scheduling, search theory)

6.6.3 Correlation of NASA CS&T and Technology Requirements

Thus far, NASA's anticipated CS&T requirements have been presented, together with an outline of relevant CS&T disciplines. In this section, they are correlated through a matrix. Each element of the matrix in table 6.7 can assume one of five values, assigned on a subjective basis by CS&T assessment team members after consultation and thorough consideration.

TABLE 6.7.— NASA COMPUTER SCIENCE AND TECHNOLOGY REQUIREMENTS

NASA — relevant computer science and technology disciplines	NASA technology requirements										
	Robotics	Smart sensors	Mission operations	Computer systems	Software	Data management	Data base systems	Management services	Human/machine services	Engineering	System engineering
(3) Applications											
(3.1) Natural Sciences	1	3	3						2		0
(3.2) Engineering	2	0	0						2	2	0
(3.4) Humanities (natural languages)			2	1							0
(3.5) Management											
(3.6) Artificial intelligence	2	3	3		1	2		0	2	1	0
(3.7) Information retrieval						2	2	1	1	0	0
(3.8) Real-time systems	3		3	2		2		0	1	1	0
(4) Software											
(4.2) Programming languages			0	0	1						
(4.3) Supervisory systems	2		1	2	2	2	2	0	0	0	
(4.4) Utility programs			1	1	1	1	1			1	0
(4.6) S/W evaluation tests and measurements	1		1		2	0	0			0	1
(5) Mathematics of computation											
(5.1) Numerical analysis	0		0		0						
(5.2) Metatheory	0			1	1						
(5.3) Combinatorial and discrete math			0	0	0					0	
(5.4) Mathematical programming	0		0	0	0						0
(5.5) Mathematical statistical problems	0	1	0	0	0						1
(5.6) Information theory			1	0	0						
(5.7) Symbolic algebraic computation	0	1	0		0						
(6) Hardware											
(6.1) Logic design/switching theory	1	1		1						0	0
(6.2) Computer systems	0	0	0	2							0
(6.3) Components and circuits	1	1		1						1	1
(8) Functions											
(8.1) Simulation and modeling	2	2	2	2	2	2	2		2	2	2
(8.2) Graphics			1	0	0	1	1		1	1	0
(8.3) O.R./decision tasks	1		1		1				1	1	0

0 = Monitor technology status and transfer technology into NASA programs.
1 = Support outside research and development.
2 = Perform in-house research, maintain peer level in state-of-the-art community.
3 = Become technology leader.

The five values are 0, 1, 2, 3, and null. A null is used where a given CS&T discipline is not expected to contribute significantly to the satisfaction of a given class of NASA technical requirements. A numerical value indicates that the CS&T discipline is expected to be a significant element in satisfying the NASA requirements in a particular area.

Further resolution is given, addressing the level of NASA commitment required to apply the CS&T disciplines successfully to the NASA requirements. Zero implies that the discipline is receiving adequate and NASA-relevant support from other sources, and the agency need only monitor the status of the technology and apply it to NASA requirements. A value of 1 means that some agency support is required to adopt a technology to applications within NASA, but R&D activities are strictly applied and can be performed through well-defined contract activities. A value of 2 implies that a substantial commitment is required to develop a discipline and apply it to NASA requirements. This commitment will involve both basic and applied research, and will establish the agency as a peer in the community of state-of-the-art researchers in the given discipline. This is a substantial commitment by NASA to a particular discipline, and will require the development of a "critical mass" of capable personnel and a stable funding environment over a period of several years. The final matrix value notation is 3, which is used in those special instances where NASA should become the recognized technology leader in a given CS&T discipline.

The correlation between NASA technology requirements of section 6.6.1 and the CS&T disciplines of section 6.6.2 are shown in matrix form in table 6.7. In general, the table shows that the agency has a wide multidisciplinary dependence on CS&T. This suggests a position of leadership for NASA in the areas of natural sciences and artificial intelligence as applied to mission operations and remote sensing, and real-time systems for robotics and mission operations. It further argues for a substantial commitment to engineering applications such as CAD/CAM technology, natural language processing, artificial intelligence and real-time systems in general, information retrieval, supervisory software, computer systems technology, simulations and modeling.

A cursory and admittedly incomplete review of existing capability within NASA suggests that state-of-the-art technology already is a part of Agency programs in the natural sciences, engineering, and simulation and modeling. Further, some good work is being done in an attempt to bring NASA's capability up to the state-of-the-art in natural language processing, although primarily through contracted research activities. But in order to fully realize the potential of CS&T within the space Agency, it appears that a substantial commitment to research in machine intelligence, real-time systems, information retrieval, supervisory systems, and computer systems is required. In many cases it was concluded that NASA has much of the requisite in-house expertise in isolated individuals and organizations, but that

the agency as a whole has been reluctant or disinterested in applying this expertise. An apparent lack of expertise *does* exist in the field of "mathematics of computation" (with a possible exception in the engineering area). This discipline can easily be overlooked as seemingly irrelevant, but in fact is a fundamental theoretical component of a broad-based and effective machine-intelligence institutional capability.

6.6.4 Facilities

To develop an institutional state-of-the-art capability in CS&T as described above will require good people and good facilities. Neither can do the job alone. Unfortunately, competent CS&T research-oriented professionals currently are in very short supply, and those few that exist are being attracted to industry and the universities through incentives of high salaries, outstanding working conditions, and intellectual freedom. None of these are offered by NASA at present, so the agency would probably be frustrated even if it were to attempt to hire the right talent. There is little NASA can do regarding salaries, so its focus in providing competitive incentives must be elsewhere. In this section, several specific recommendations are made with respect to facilities which the CS&T team considers prerequisite to any serious attempt by NASA to develop significant in-house capabilities in CS&T.

Interactive, on-line programming environment. Most programming is currently done within NASA on 10-year-old batch-oriented computer systems, where programmers still manipulate card decks and experience turnaround times measured in hours or even days. In order to attract competent researchers and to provide an environment in which they can labor productively, an absolute prerequisite is a fully interactive, on-line programming environment. For instance, Teitelman (1979) describes a typical state-of-the-art interactive system of the type required. NASA will find that this type of system, when made generally available to its personnel, will yield a very significant increase in programmer productivity. It is expected that this increase will be sizable enough to more than offset the additional cost of the on-line capability.

NASA has historically met its computing requirements through the purchase of computing equipment (e.g., instead of leasing). Due to the intricacies of the government ADP procurement process, 5 years typically will lapse between the conception of a new system and its actual operation — and then that system will remain in operation for 10-15 years, so that a system will be 15-20 years behind the state-of-the-art at its retirement (and 10-15 years behind during the "prime" of its life). NASA may wish to consider as an alternative for its nonmission (and, specifically, R&D) computing requirements the purchase of time-sharing services from a quality commercial vendor, so that it always has access to the best of the commercial offerings at any given time.

Computer communications network. NASA currently lacks any effective mechanism to provide digital communication between its computers in a general way. Several small efforts at individual centers have addressed intercomputer communication, and NASA has actively participated in international negotiations on intercomputer communication protocols, but no agency-wide effort has been made to apply this technology within NASA.

A strong case can be made for NASA to develop such a capability. It facilitates regular communication among geographically dispersed personnel and enables the sharing of both hardware and software resources. This can be particularly important for coordinating joint research among the centers.

One is tempted to envision within NASA a network structured logically as a hierarchy, in which the "standard equipment" (much as is a desk and chair) of an individual is a terminal which gives him direct access to local computing resources. These local computers are then aggregated into local computer networks to provide load balancing and resource sharing for a community of users as well as access to resources outside the local network such as other local networks within NASA, and extending to include the ARPANET and commercial facilities such as TYMNET and TELENET. It is not difficult to consider a NASA network which links all individuals within a Center, Centers to each other and to Headquarters, and provides access to non-NASA resources. There is little doubt that such a system will eventually become a reality. The telephone system currently provides just this type of capability for voice communication. The question becomes more one of *when* than *if* it will happen.

The ARPANET was developed largely as an experiment in this type of technology and has evolved into a primary vehicle for communication among researchers in the artificial intelligence community as well as other CS&T disciplines. As a first step toward integrating this type of capability into NASA systems, the agency should seriously consider negotiating with the Defense Communications Agency (DCA) for agency-wide access to the ARPANET. Not only would such a step provide a communications link with a large part of the research community of CS&T, but it would also provide the opportunity to perform communications experiments within NASA with a minimum investment of agency resources. This experience should equip NASA with the knowledge and expertise it will need to consider the implementation of an in-house networking capability.

Office automation. A significant proportion of NASA's resources are consumed manipulating documentation in many forms, including standard government forms, design documentation for software and hardware systems, intercenter and inter-agency agreements, and scientific papers; yet most of these processes are largely performed through manual means. Given on-line, interactive systems and a good communications capability, it becomes a minor step to provide a word processing capability which enables the author of a document to generate a document, have it reviewed by others, revise the document, and publish it without typing any portion of it more than once, and without standing over a copying machine and circulating review copies.

The current state of the art in "expert systems" is well suited to managing standardized administrative forms. One can easily envision a "Travel Expert System," for example, which knows government travel regulations, per diem rates, etc., and could interactively assist an individual (secretarial or professional) in constructing a set of travel orders. Such a system would also know the approval required and could automatically route the orders to the appropriate signature authorities. Changes in travel regulations would then be integrated into the expert system directly and applied as necessary, avoiding the costly notification process to all concerned individuals, with the assurance that everyone will be using up-to-date information.

Utilizing state-of-the-art office automation technology within NASA to manage documentation, coordinate manpower, and provide communication among personnel could significantly improve the productivity of NASA personnel. Integrating state-of-the-art machine intelligence capabilities into such an office environment could provide untold improvements in the efficiency and effectiveness of the organization, including the potential for significantly enhanced project management and rigorous statistical exploration of policy alternatives and their impacts.

6.6.5 Organizations

NASA is not presently organized to support a vigorous program in CS&T. The most apparent lack is a discipline office at the Headquarters level which supports research and development in computer science and serves as an agency advocate for the incorporation of state-of-the-art capabilities into NASA programs. There presently exist within the space agency many computer scientists capable of pursuing state-of-the-art research and of integrating contemporary technology into NASA programs, but there is no place for them to go for support other than mission-oriented offices whose goals and objectives are not consistent with supporting long-term commitments in CS&T R&D.

In addition to recognizing the requirement for a Headquarters CS&T discipline office, the study team fully supports the recommendation of the NASA Study Group on Machine Intelligence and Robotics that an advisory council composed of industry leaders in CS&T should be formed. This council would assist the agency in developing its computer science programs in order to assure a proper focus and to construct the appropriate relationships with other research organizations.

It is beyond the scope of the present study to recommend how NASA should organize institutionally to develop its CS&T capabilities, but several ideas have surfaced which may be useful to the agency in its consideration of future courses of action. In order to be effective, it would seem appropriate that NASA's CS&T endeavors maintain a multi-mission focus. A possible starting point may be a nucleus of discipline specialists to develop the program coupled with an agency-wide matrix management strategy to apply contemporary CS&T in mission environments and a vigorous encouragement of the development of CS&T "centers of excellence" at the Centers. Before embarking on any major organizational changes, however, it is useful to perform a systems analysis to fully explore the organizational possibilities and their ramifications. In this regard, the techniques developed by Krone (1980) are highly recommended.

In consideration of the two-fold objective of maintaining state-of-the-art expertise in CS&T and applying this expertise to NASA programs, one is confronted with the dilemma of providing both an effective R&D environment and a line organization capability to apply CS&T to real missions. If one assumes the existence of line organizational entities that now exist within NASA but applied to CS&T endeavors, then a possibility to be considered is the augmentation of the line management positions with staff researchers as illustrated in table 6.8. The positions of "Fellow" are intended to be highly competitive and attractive positions open to employees of NASA or other government agencies, industry, and academia. They might be treated similarly to professorships within universities, where highly talented individuals may be tenured in a position, but many are rotated through positions on temporary assignments of several years' duration.

An intriguing organizational structure apparently has been developed by the Navy for its new artificial intelligence laboratory at the Naval Research Laboratory (NRL)

in Washington, D.C. A rough outline of the structure is shown in figure 6.3. The major point of interest regarding the NRL effort is that it is organizationally constructed to maximize scientific productivity. The head of the organization is the Chief Scientist, who is expected to contribute in a meaningful scientific way to the work of the organization. The CS&T study group was able to find out rather few details regarding the proposed operation of the NRL facility, but it appears to be sufficiently interesting that NASA may wish to explore this alternative during the system analysis phase.

6.6.6 Programs for Excellence

Perhaps the most fundamental requirement in maintaining one's technical excellence is to maintain active relationships with peer researchers. This will involve both formal and informal interfaces with standards organizations, other government agencies, universities, industrial R&D programs, and professional societies. A good set of computational facilities and communications capabilities as proposed in section 6.6.4 will facilitate this process.

Participation in joint government/industry/academic programs such as institutes and consortiums can formally provide not only a mechanism for applying more leverage to technical problems but also, potentially, a very appropriate forum for technical interchange. "Visiting Scientist" programs, where NASA sends selected individuals to major research environments such as MIT, SRI, and Xerox/PARC for periods of 6 months to a year can be very effective in transferring state-of-the-art concepts and technology into NASA programs. The Agency may also wish to consider Scientist Exchange Programs, in which NASA scientists perform research in university or industrial environments while their counterparts work at NASA for 6 months to a year. In some cases, close relationships between field centers and local universities may be mutually beneficial, and

TABLE 6.8.— POSSIBLE PERSONNEL POSITIONS
FOR CS&T RESEARCH

Pay, status	Line management	Applied research and consulting	Independent research
	Division Chief	Senior Staff Scientist	Senior Research Fellow
	Branch Head	Staff Scientist	Associate Research Fellow
	Section Head	Research Scientist	Assistant Research Fellow
		Responsibility	Freedom
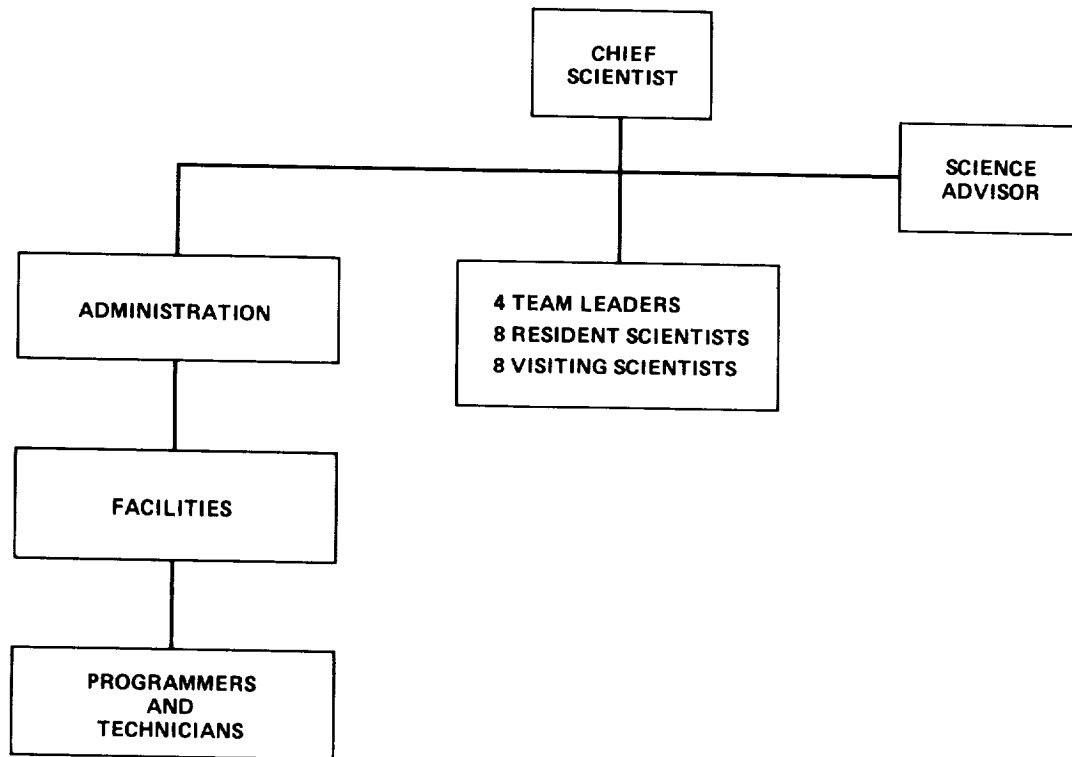
Figure 6.3.— NRL Artificial Intelligence Laboratory organization.

could include adjunct professorships as well as sponsoring graduate student thesis research at NASA facilities. This could prove to be an effective recruiting device. The agency may also wish to consider student loan programs for graduate students, wherein part of the loan is forgiven if the student completes an advanced degree and comes to work for NASA. A final suggestion is to sponsor Ph.D. thesis competitions, in which a "Space Technology Award," say, of perhaps $5000 is awarded annually to the best thesis relating to problems relevant to NASA. In general, the CS&T study group believes that there are many institutional programs which will cost NASA very little, yet can do much to maintain a NASA capability once it is obtained. Within NASA are examples of how programs such as these have succeeded particularly well in physics and the space sciences.

6.6.7 Summary and Conclusions

The primary technical components of a NASA program in Computer Science and Technology include research in

machine intelligence, real-time systems, information retrieval, supervisory systems, and computer systems. NASA's computing and communications facilities require substantial upgrading in order to perform the proposed research and to attract competent personnel. An interactive, on-line programming environment appears essential, as well as a move by NASA to provide extensive communication capabilities among its computers and evolve toward an agency-wide computer network. It is further recommended that NASA consider the potentials of office automation technology for routine administrative work and document management, and explore the utility of machine intelligence in project management and policy analysis.

The organizational structure required to perform both state-of-the-art research and to apply modern CS&T was considered. The team concluded that this topic deserves a complete organizational analysis of alternatives, a task which can most effectively be done within NASA itself. Such a study should be given high priority in consideration of responses to requirements for implementing an advanced machine intelligence based program within NASA.

6.7 References

Amarel, S.: On Representations of Problems of Reasoning About Actions. In Machine Intelligence, D. Michie, ed., vol. 3, Edinburg Univ. Press, Edinburg, 1968.

ASC, "NASA Data/Information System-Related Technology Trends," prepared for NASA Goddard Space Flight Center by the Analytic Sciences Corporation (ASC) under contract NAS5-25739, May 1980.

Balzer, Robert et al., eds.: Proceedings of the First Annual National Conference on Artificial Intelligence, Stanford University, 18-21 August 1980. American Association of Artificial Intelligence and Stanford University, Menlo Park, Garcia-Robinson, 1980.

Baum, Michael: Giving a Robot the Eye, Dimensions, vol. 63, no. 4, April 1979, p. 3.

Bejczy, Antal K.: Advanced Teleoperators. Astronautics and Aeronautics, vol. 17, May 1979, pp. 20-31.

Bobrow, D. G.; and Winograd, T.: Experience with KRL-O: One Cycle of a Knowledge Representation Language. Proc. of the Fifth International Joint Conference on Artificial Intelligence, MIT, Cambridge, Massachusetts, 1977, pp. 213-222.

Bradley, William E.: Telefactor Control of Space Operations. Astronautics and Aeronautics, vol. 5, May 1967, pp. 32-38.

Bradley, William E.: Immediate Program to Provide Telefactor Systems for Demonstration and Test in NASA Centers. Personal communication, 5 August 1980.

Breckenridge, R. A.; and Husson, C.: Smart Sensors in Spacecraft: Impact and Trends. Progress in Astronautics and Aeronautics, vol. 67, AIAA, New York, 1979, pp. 3-12.

Brown, J. S.; and Burton, R. R.: Multiple Representations of Knowledge for Tutorial Reasoning. In Representations and Understanding, Daniel G. Bobrow and Allan Collins, eds., Studies in Cognitive Science. Academic Press, New York, 1975, pp. 311-349.

Buchanan, B. G.; and Lederberg, Joshua: The Heuristic DENDRAL Program for Explaining Empirical Data. IFIP Congress, 71, Ljubljana, Yugoslavia, Aug. 23-28, 1971. Information Processing 71, Proceedings. C. V. Freiman, Editor. Amsterdam, North-Holland Publ. Co., 1972, pp. 179-188.

Chin, C. H., ed.: Joint Workshop on Pattern Recognition and Artificial Intelligence. Institute of Electrical and Electronics Engineers, June 1-6, 1976, Hyannis, Mass. 1976.

Computing Reviews: Categories of the Computing Sciences, vol. 17, no. 5, May 1976, pp. 172-195.

Corliss, William R.; and Johnsen, Edwin G.: Teleoperator Control, NASA SP-5070, 1968.

Duda, Richard O.; and Hart, Peter E.: Pattern Classification and Scene Analysis. John Wiley & Sons, New York, 1973.

Duda, R. O. et al.: Development of the PROSPECTOR Consultation System for Mineral Exploration. Final Report, SRI International, Projects 5821 and 6415, Menlo Park, Calif., October 1978.

EER: Information Science Program Plan for NASA/ Goddard Space Flight Center. Prepared by Engineering and Economics Research, Inc. (EER), under contract NAS5-73942-B, June 1980.

Erman, L. D.; and Lesser, V. R.: A Multi-Level Organization for Problem Solving Using Many, Diverse, Cooperating Sources of Knowledge. Proceedings of the Fourth International Joint Conference on Artificial Intelligence, Tbilisi, Georgia, USSR, Sept. 1975, pp. 483-490.

Ernst, George W.; and Newell, Allen: GPS: A Case Study in Generality and Problem Solving. Academic Press, New York, 1969.

Fann, K. T.: Peirce's Theory of Abduction. Martinus Nijhoff, The Hague, 1970.

Feigenbaum, E. A.: The Art of Artificial Intelligence — Themes and Case Studies of Knowledge Engineering. Proc. Fifth International Joint Conference on Artificial Intelligence, MIT, Cambridge, Massachusetts, 1977.

Feigenbaum, E. A.; Buchanan, B. G.; and Lederberg, J.: On Generality and Problem Solving: A Case Study Using the DENDRAL Program in Machine Intelligence, B. Meltzer, D. Michie, eds., vol. 6, Halsted Press, New York, 1971, pp. 165-190.

Fikes, R. E.; and Nilsson, N. J.: STRIPS: A New Approach to the Application of Theorem Proving to Problem Solving. Artificial Intelligence, vol. 2, 1971, pp. 189-208.

Gilmore, P. A.; Batcher, K. E.; Davis, M. H.; Lott, R. W.; and Buckley, J. T.: Massively Parallel Processor. Phase I Final Report, prepared by Goodyear Aerospace Corporation for Goddard Space Flight Center, July 1979.

Goldstein, I. P.; and Roberts, R. B.: Nudge: A Knowledge-Based Scheduling Program. Proc. of the Fifth International Joint Conference on Artificial Intelligence, MIT, Cambridge, Massachusetts, 1977, pp. 257-263.

Goulomb, S.; and Baumert, L.: Backtrack Programming. J. ACM, vol. 12, 1965, pp. 516-524.

Hajeck, P.; and Havranek, T.: 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 Univ. Press, Cambridge, 1958.

Haralick, R. M.: Texture-Tone Study with Application to Digitized Imagery. Final Tech. Summary, ETL-CR-74-17, Contract DADKO2-70-C-03888, August 1974.

Haye, R. M.: Smart Remote Sensor Needs for U.S. Coast Guard Ocean Environmental Missions. Progress in Astronautics and Aeronautics, vol. 67, AIAA, New York, 1979, pp. 44-59.

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 Univ., July 1980.

Heer, Ewald: Toward Automated Operations in Space: An Introduction. Astronautics and Aeronautics, vol. 17, May 1979, pp. 16-18.

Hewitt, C.: Description and Theoretical Analysis (Using Schemata) of PLANNER: A Language for Proving Theorems and Manipulating Models in a Robot. Artificial Intelligence Laboratory Report AI-TR-258, MIT, Cambridge, Massachusetts, 1972.

Ho, Y. C.: Team Decision Theory and Information Structures. Proc. of the IEEE, vol. 68, June 1980, pp. 644-654.

IEEE Proceedings on Pattern Recognition and Image Processing, May 1979.

Kalush, R. J., Jr.: The Problem of Resolution in the LANDSAT Imagery. Remote Sensing Quarterly, vol. 2, 1980, pp. 3-48.

Kowalski, R. A.: Predicate Logic as Programming Language. Proc. IFIP 1974, North-Holland, Amsterdam, 1974, pp. 569-574.

Kraiss, F.: Decision-Making and Problem Solving with Computer Assistance. NASA TM-76008, January 1980.

Krone, R. M.: Systems Analysis and Policy Sciences, Theory and Practice. John Wiley & Sons, Inc., 1980.

Kuipers, B. J.: Representing Knowledge of Large-Scale Space. Artificial Intelligence Laboratory Report AI-TR-418, MIT, Cambridge, Massachusetts, 1977.

Lakatos, Imre: Falsification and the Methodology of Scientific Research Programs. In Criticism and the Growth of Knowledge, Imre Lakatos and Alan Musgrave, eds., Univ. Press, Cambridge, 1970a, pp. 91-196.

Lakatos, Imre: The Changing Logic of Scientific Discovery. University Press, Cambridge, 1970b.

Lakatos, Imre: Proofs and Refutations: The Logic of Mathematical Discovery. Cambridge Univ. Press, Cambridge, 1976.

Lighthill, Sir James: Artificial Intelligence: A Paper Symposium, Science Research Council, Great Britain, 1972.

Lippay, Andrew L.: Multi-Axis Hand Controller for the Shuttle Remote Manipulator System. Paper delivered at the 13th Annual Conference on Manual Control, June 1977.

Martin, W. A.; and Fateman, R. J.: The MACSYMA System. In Proc. of the ACM Second Symposium on Symbolic and Algebraic Manipulation, S. R. Petrick, ed., Los Angeles, Calif., March 1971, pp. 23-25.

Matsushima, H.; Uno, T.; and Ejur, M.: Image Processing by Experimental Arrayed Processor. 6th International Joint Conference on Artificial Intelligence, Tokyo, 1979, pp. S13-S15.

Meade, E.; and Nedwich, B.: Manipulator Development Facility Description. NASA JSC-11029, Rev. A, January 1978.

Miller, R. H.; and Smith, David B. S.: Extraterrestrial Processing and Manufacturing of Large Space Systems, Volumes 1-3, NASA Contract NAS8-32925, NASA CR-161293, September 1979.

Minsky, M.: A Framework for Representing Knowledge. In The Psychology of Computer Vision, P. H. Winston, ed., McGraw-Hill, New York, 1975, pp. 211-277.

Minsky, M.: Toward a Remotely-Manned Energy and Production Economy. MIT Artificial Intelligence Laboratory, A.I. Memo No. 544, September 1979. 19 pp. See also "Telepresence," Omni, vol. 2, June 1980, pp. 44-52.

Mitchell, O. R.; Meyers, C. R.; and Boyne, W. A.: A Max-Min Measure for Image Texture Analysis. IEEE Transactions on Computers, vol. C-25, 1977, pp. 408-414.

Morgan, C. S.: Hypothesis Generation by Machine. Artificial Intelligence, vol. 2, 1971, pp. 179-187.

Morgan, C. S.: On the Algorithmic Generation of Hypothesis, Scientia, vol. 108, 1973, pp. 585-598.

Murphy, L. P.; and Jarman, J. W.: The Role of Smart Sensors in Earth Resources Remote Sensing Programs. Progress in Astronautics and Aeronautics, vol. 67, AIAA, New York, 1979, pp. 101-109.

NASA Advisory Council: Future Directions for the Life Sciences in NASA. Report of the Life Sciences Advisory Committee of the NASA Advisory Council, November 1978.

NASA SP-387: A Forecast of Space Technology, 1980-2000, 1976.

Newell, A.; and Simon, H. A.: GPS: A Program that Simulates Human Thought. In Computers and Thought, E. A. Feigenbaum, J. Feldman, eds., McGraw-Hill, New York, 1963, pp. 279-293.

Nilsson, Nils J.: Problem-Solving Methods in Artificial Intelligence. McGraw-Hill, New York, 1971.

Nilsson, Nils J.: Artificial Intelligence. Information Processing 74, vol. 4, Proceedings of IFIP Congress 74, 1974, pp. 778-801.

Nilsson, Nils J.: Principles of Artificial Intelligence. Tioga, Palo Alto, California, 1980.

Office of Aeronautics and Space Technology (OAST): NASA Space Systems Technology Model, volumes 1-3, NASW-2937, NASA Headquarters, Washington, D.C., May 1980.

Peirce, Charles Sanders: Collected Papers, vols. 1-6, Charles Hartshorne and Paul Weiss, eds. The Belknap Press of Harvard Univ. Press, Cambridge, 1960.

Peirce, Charles Sanders: Collected Papers, vols. 7-8. Arthur W. Burks, ed. The Belknap Press of Harvard Univ. Press, Cambridge, 1966.

Pohl, I.: Practical and Theoretical Considerations in Heuristic Search Algorithms. In Machine Intelligence, vol. 8, E. W. Elcock, D. Michie, eds., Ellis Horwood Ltd. (Halsted Press), 1977, pp. 55-72.

Pople, H. E., Jr.: The Formation of Composite Hypotheses in Diagnostic Problem-Solving: An Exercise in Synthetic

Reasoning. Proc. of the Fifth International Joint Conference on Artificial Intelligence, MIT, Cambridge, Massachusetts, 1977, pp. 1030-1037.

Radner, R.: Team Decision Problems. Ann. Math. Stat., vol. 33, 1962, pp. 857-881.

Raibert, M. H.: Autonomous Mechanical Assembly on the Space Shuttle: An Overview. NASA CR-158818, July 1979.

Reiter, R.: The Use of Models in Automatic Theorem-Proving. Technical Report 72-09, Dept. of Computer Science, Univ. of British Columbia, Canada, 1972.

Reiter, R.: A Logic for Default Reasoning. Artificial Intelligence, vol. 13, 1980, pp. 81-132.

Rich, E.: Building and Exploiting User Models. Proc. of 6th International Joint Conference on Artificial Intelligence. Tokyo, 1979, pp. 720-722.

Sacerdoti, E. D.: Planning in a Hierarchy of Abstraction Spaces. Artificial Intelligence, vol. 5, 1974, pp. 115-135.

Sacerdoti, E. D.: Problem Solving Tactics. Proc. Sixth International Joint Conference on Artificial Intelligence. Tokyo, Japan, 1979, pp. 1077-1085.

Sagan, Carl, chmn.: Machine Intelligence and Robotics: Report of the NASA Study Group. NASA-JPL Report No. 715-32, March 1980.

Sandell, N. R.; Varaiya, P.; Athans, M.; and Safonov, M. G.: Survey of Decentralized Control Methods for Large Scale Systems. IEEE Trans. Automat. Contr., vol. AC-23, April 1978, pp. 108-128.

Sandford, D. M.: Using Sophisticated Models in Resolution Theorem Proving. Lecture Notes in Computer Science, vol. 90, Springer-Verlag, New York, 1980.

Schappell, R. T.: A Study of Early Intrusion into Algorithm Development Utilizing Information Processing Advanced Technology. Martin Marietta Corp., P-80-48052-1, 1980a.

Schappell, R. T.: FILE Technology Developments for Remote Sensing. SPIE Annual Int. and Tech. Symposium and Exhibit, San Diego, California, 31 July 1980b.

Schappell, R. T.; Polhemus, J. T.; Lowrie, J. W.; Hughes, C. A.; Stephens, J. R.; and Chang, Chieng-Y.: Applications of Advanced Technology to Space Automation. NASA CR-158350, January 1979a.

Schappell, R. T.; and Tietz, J. C.: Landmark Identification and Tracking Experiments. Progress in Astronautics and Aeronautics, vol. 67, AIAA, New York, 1979, pp. 134-151.

Schappell, R. T.; Vandenberg, F. A.; and Hughes, C. A.: Study of Automated Rendezvous and Docking Technology: Final Report. Martin Marietta Corporation, 1979b.

Schlienn, S.: Geometric Correction Registration and Resampling of LANDSAT Imagery. Canad. J. Remote Sensing, vol. 5, 1979, pp. 74-89.

Shin, C. N.; and Yerazunis, S.: Data Acquisition and Path Selection Decision-Making for an Autonomous Roving Vehicle. Progress Report, Rensselaer Polytechnic Institute, TR MP-62, 1978.

Shortliffe, E. H.: Computer-Based Medical Consultations: MYCIN. Elsevier, New York, 1976.

Space Shuttle, JSC, 1976: Space Transportation User Handbook, JSC, July 1977; Space Shuttle System Payload Accommodations, JSC-07700, vol. XIV, Revision F, 22 September 1978.

Spann, G. W.: Satellite Remote Sensing Markets in the 1980's. Photogrammetric Engineering and Remote Sensing, vol. 46, 1980, pp. 65-69.

Srinivasan, C. V.: The Architecture of Coherent Information Systems: A General Problem Solving System. IEEE Trans. on Computers, vol. C-25, 1976, pp. 390-402.

Srinivasan, C. V.: Knowledge-Based Learning, An Example. DCS-TR-90, paper presented at Workshop on Machine Learning, Carnegie-Mellon Univ., July 1980.

Srinivasan, C. V.; and Sandford, D. M.: Knowledge-Based Learning Systems, Part I: A Proposal for Research; Part II: An Introduction to the Meta-Theory; Part III: Logical Foundations. DCS-TR-89, Dept. of Computer Science, Hill Center, Busch Campus, Rutgers Univ., New Brunswick, New Jersey, March 1980.

Sussman, G. J.: A Computer Model of Skill Acquisition. American Elsevier, New York, 1975.

Teitelman, W.: A Display Oriented Programmer's Assistant. Int. J. Man-Machine Studies, November 1979, pp. 157-187.

Tenenbaum, Jay M.: Reconstructing Smooth Surfaces. In Proc. Image Understanding Workshop, Univ. of Southern California, November 1979.

Thorley, G. A.; and Robinove, C. J.: Current and Potential Uses of Aerospace Technology by the U.S. Department of the Interior. Progress in Astronautics and Aeronautics, vol. 67, AIAA, New York, 1979, pp. 15-26.

Vahey, D. W.: Smart Remote Holographic Processor Based on the Materials Characteristics of $LiNbO_3$. Progress in Astronautics and Aeronautics, vol. 67, AIAA, New York, 1979, pp. 352-366.

Van de Brug, G.; and Minker J.: State-Space, Problem-Reduction, and Theorem Proving – Some Relationships. J. ACM, vol. 18, 1975, pp. 107-115.

Van Emden, M.: Programming with Resolution Logic. In Machine Intelligence, E. W. Elcock and D. Michie, eds., Ellis Horwood Ltd. (Halsted Press), vol. 8, 1977, pp. 266-299.

Warren, D.; and Pereira, L. M.: PROLOG – The Language and its Implementation Compared with LISP. Proc. Symposium on Artificial Intelligence and Programming Languages, ACM, 1977.

Whitney, M. W.: Processing and Storing Information. In A Forecast of Space Technology, 1980-2000. NASA SP-387, section IV, 1976.

Williams, O. W.: Outlook on Future Mapping, Charting and Geodesy Systems. Photogrammetric Engineering and Remote Sensing, vol. 46, 1980, pp. 487–490.

Wilson, G. A.; and Minker, J.: Resolution, Refinements, and Search Strategies: A Comparative Study. IEEE Transactions on Computers, vol. C-25, 1976, pp. 782–801.

Winston, Patrick H.: Progress in Artificial Intelligence. Vols. I and II, 1978.

Woods, W. A.: What's in a Link: Foundations for Semantic Networks. In Representation and Understanding, D. G. Bobrow and A. Collins, eds., Academic Press, New York, 1975, pp. 35–82.

Woods, W. A.; Kaplan, R. M.; and Nash-Webber, B.: The Lunar Sciences Natural Language Information System: Final Report, BBN Report No. 2378, Bolt, Beranek and Newman, Cambridge, Massachusetts, 1972.