A ROBOTICS PERSPECTIVE ON HUMAN SPACEFLIGHT

ALEX ELLERY
School of Engineering, Kingston University, London SW15 3DW, UK
E-mail: Ku17681@atlas.kingston.ac.uk

Abstract. There has traditionally been a dichotomy in the space community regarding the efficacy of human versus robotic exploration of space. I argue that no such dichotomy is necessary, and that there is a natural and synergistic division of labour between man and machine, and that this division of labour will evolve in symbiotic fashion. The present state-of-the-art robotics technology is insufficient to replace the human in space, but is sufficient to act as a useful, even necessary, tool in aiding the astronaut in the conduct of useful work. I further argue that as robotics technology advances, the human will be further relieved to perform tasks best suited to human decision-making and flexibility that is unlikely in the near-term to be matched by autonomous or teleoperated machines.

Keywords: Autonomy, exploration, in-orbit, manned, Mars, planetary, robotic, space, telerobotic, unmanned

1. Introduction

The division in the scientific community that polarises scientists into opposing camps regarding manned spaceflight has a history as long as the space race itself. The debate over humans versus robots for space exploration has often been seen as a mutually exclusive one, particularly from the perspective of funding allocation being a zero-sum game. The detractors of manned spaceflight maintain that human spaceflight is wasteful of financial resources on the basis that it does not provide adequate return on investment, the return being defined either financially or scientifically. The primary reason for this is that scientific return takes a backseat to other more public relations issues regarding men in space. The proponents of manned spaceflight argue however that the scientific returns can be greatly in excess of the capabilities of robotic approaches. The most often quoted example is that vastly greater amounts and variety of lunar rock from a greater range of intelligently-selected terrain (382 kg) were recovered by the Apollo astronauts compared to the Russian Lunakhod sample return missions (321 g). Our scientific understanding of the geological history of the moon is almost entirely attributable to those Apollo samples (Zorpette, 2000). There is little doubt that manned spaceflight affords a flexibility of space operations that is unattainable in unmanned, robotic missions. However, robotic missions offer greater range than manned missions allowing access to the outer solar system. I argue that such polarised disputes are simplistic and take no account of the capabilities and limitations of robotic auto-

Earth, Moon and Planets **87:** 173–190, 2001.

mation in spacecraft and that manned and robotic spaceflight are complementary and indeed symbiotic in nature.

Unmanned robotic systems are essential for the success of manned missions, especially planetary exploration. Robotic precursor missions are essential for initial surveys and characterisation of planetary landing sites for human missions. Furthermore, robotics capabilities in Earth orbit would considerably enhance the missions of in-orbit spacecraft, and there is no doubt that the potential for robotic systems to support manned space activities are great (Ellery, 2000). It is worth noting that the dichotomy between robotics and human spaceflight have not always been seen to be so irreconcilable. During the early days of the space station programme, robotics as a space technology was seen as an important aspect of space infrastructure development for the support of manned space activities. In 1985, the NASA Advanced Technology Advisory Committee recommended to the US Congress Committee on Science and Technology that no less than 10% of the then manned space station programme budget be devoted to robotics and automation (Cohen and Erickson, 1985). However, the flagships of that commitment, the Orbital Maneouvring Vehicle and the Flight Telerobotic Servicer were the first to feel the axe under budget constraints.

I argue here that the nature of robotics technology lends itself to certain tasks in space or on planetary bodies that are best suited to machines and that there are other essential tasks that must be performed by human beings in space and on planetary surfaces. By considering the limitations on robotics technology, we can devise an optimal division of labour between man and machine in the exploration of the space environment.

2. The Hierarchical Approach to Telerobotics

The robot is a machine that is designed ultimately to reproduce or exceed human capabilities in potentially hazardous, remote environments such as space. Robotics as a discipline is primarily concerned with control systems which essentially derive their behaviour on the basis of their sensory data. This essentially means feedback of sensory data in order to modify the robot's actions on the environment. The robotics problem *is* the control problem. Robotic control systems are defined by a hierarchy of different levels defined by their temporal characteristics. At the most primitive level, a robotic control system is implemented by the servo-control loop. Sensory feedback must be rapid to track changes in the environment that would otherwise lead to erroneous actions on the part of the robot. Response times are from milliseconds to around one-tenth of a second – this is the level at which lies teleoperation. It is characterised typically by numerical computations. At the most sophisticated level, the robotic control system is implemented with artificial intelligence programs whereby plans are derived as solutions to problems with much slower feedback times of the order of minutes or hours – this is the level of a

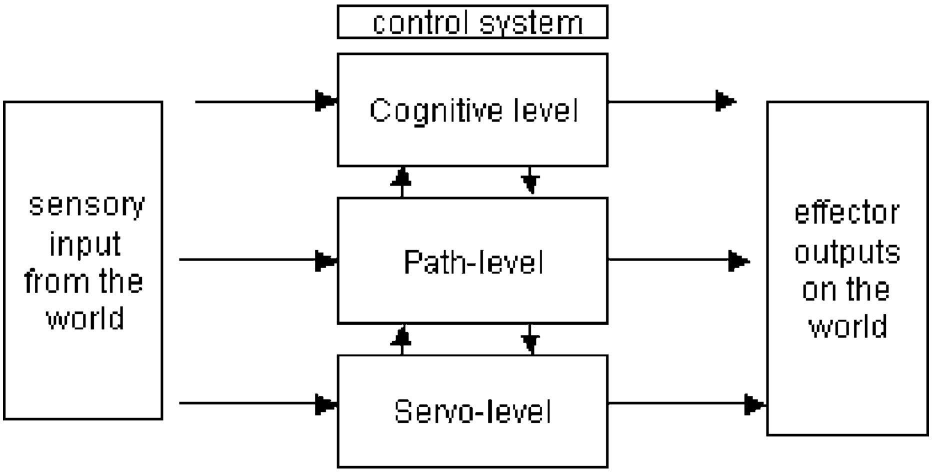

Figure 1. Three-layer control hierarchy.

smart telerobot capable of on-line decision-making. This level is characterised by symbolic computations. In between these levels lie a mixture of both. For instance, in any form of controlled motion, be it a rover or a manipulator, this level is typified by path control and obstacle avoidance. Obstacle avoidance is a major concern in robotics as it relies on real-time sensory data, yet it affects higher level plans which may require adjustment or even re-planning. The characteristic timescales involved depend on the speed of actuation and the rapidity of change in the environment. Neural networks have been applied to this level of robot control by representing path control as an adaptive pattern recognition problem which defines the mapping between sensory inputs and motor outputs, but this method took no account of obstacle avoidance (Kuperstein et al., 1987).

3. The Division of Labour in the Man–Machine Society

Consistent, repetitive and routine tasks requiring high precision, speed and repeatability are ideally suited to automation since machines do not tire or become inattentive. Most civil airline accidents are caused by human error due to limitations in vigilance or faulty judgements (Rouse and Cody, 1987). Manual tasks that can be performed prescriptively by detailed, algorithmic procedures are eminently suitable for automation. Such pre-programmed sequences have been used on spacecraft for autonomous execution – the Viking lander on Mars (1976) utilised an articulated manipulator to recover samples from the Martian soil for insertion to the onboard biochemical experiments (Schmidt and Hawes, 1999). But such simple tasks in controlled environments however have not always been carried out as expected, e.g., the failure of the Galileo (1989) high gain antenna to deploy correctly. However, automation of simple pre-programmable tasks can relieve the human being from manual workloads which in space can be arduous and dangerous. For robotic manipulators and mechanisms, these deterministic tasks are driven by geometry, e.g., peg-in-hole tasks. Machines are also ideally suited to tasks like monitoring multiple complex systems, and for fault detection and correction. This was part of the function of the Deep Space 1 probe's Remote Agent software that was developed to achieve greater autonomy in spacecraft operations. Automated fault monitoring and diagnosis can provide a valuable real-time aid to humans acting in an overseeing decision-making capacity. However, this highlights *the* major design issue in man–machine interfacing: The majority of operational errors that occur usually result from human misinterpretation of information due to information overload. The USS Vincennes mistakenly shot down the civil airliner Iran Air 655 killing all 290 passengers in 1988. This tragedy was entirely attributed to the lack of adequately designed visual man–machine interfaces implemented on the US Navy's radar equipment. "User-friendliness" is essential – humans are primarily visual creatures and graphical representation of data is generally the most appropriate approach. Humans are suited to robust pattern recognition and data interpretation in noisy and uncertain environments and for reactive responses to unexpected occurrences. They are best suited to non-repetitive, highly complex tasks characterised by high variability. Image processing algorithms are often prone to failure in natural environments which are characterised by fractal shapes, though they perform much better in environments with well-defined geometrical objects such as man-made engineered objects. These non-deterministic tasks are those which require extensive dexterity skills or extensive "thinking-on-one's feet" skills and involve processes that are not well-understood as they are subject to many changing factors. Machines cannot yet, nor will they in the near future, emulate human inductive decision-making and the utilisation of generalisations from past experience.

Telerobotics may be regarded as a specialised form of man-machine interface, but there are various modes of shared man-machine control possible, reflecting the various degrees of human supervision. Teleoperation involves the projection of the senses, vision and touch typically, to the physically remote environment where the robot resides. In this case, control is directly implemented by the human being. Effective reflection of sensory feedback of visual and touch data back to the human in-the-loop is essential for teleoperation (Bejczy, 1980). The most obvious development of this approach is in virtual reality technology which seeks to maximise the transparency between the human and the machine to the extent that the human experiences "telepresence". Automation at the other end of the scale implies that all tasks are executed autonomously without any human intervention, though some form of human supervision is generally implemented. Between these extreme cases, there are many shades of shared division of labour between man and machine with the human taking on different degrees of supervision.

4. Carry On Up Hierarchy

Teleoperation becomes difficult, inefficient or even impossible if the remote site Invokes a round-trip signal propagation time delays beyond a few seconds. Beyond 1-second time delays, move-and-wait strategies must be employed. Signal time-of-flight is not the only delay – signal processing and computer network propagation increase this time, so a time delay of up to 1.5 s or more can be incurred for a robot spacecraft in geostationary orbit at 36,000 km altitude controlled from a ground station. Even in low earth orbit, the need for continuous communication between a robot freeflyer spacecraft and the ground necessitates relaying the signal through a geosynchronous system such as TDRSS (Tracking and Data Relay Satellite System). Furthermore, delays in force feedback control between ground and remote sites can cause instability more rapidly than in visual feedback situations. In effect, direct teleoperation is not an efficient strategy if it requires tasks to be performed remotely with more than 1.0–1.5 s delay windows. Evidently, the problem is worsened in cases like lunar missions where the round-trip delay increases to 2.6 s (though efficiency may not necessarily be a major concern in this case), and Mars missions in particular where the delay time can be up to 40 minutes. Graphical predictive displays such as those based on the Smith predictor algorithm can offer limited improvement in performance for short time delays of up to a few seconds. Shared control with a mixture of automation and human supervision is the best approach – so called supervised autonomy. The well-structured nature of many phases of robotic motion such as pick-and-place, tool exchange from toolbox jigs and orbital replacement unit (ORU) module exchange operations lend themselves to automation. Indeed, many of these manipulator tasks are executed repeatedly so re-use of pre-programmed, automated subroutines can greatly relieve the workload on the human teleoperator. Illumination for instance is a considerable problem in the space environment requiring constant adjustment during teleoperation: automated control of camera and lighting has a significant effect on teleoperator performance.

The typical approach to robotic automation is to minimise the environmental complexity that the machine has to operate in. Pick-and-place robotics is the archetypal example of this – the environment is artificially designed and constrained to ensure that it is highly predictable with minimal uncertainty. Peg-in-hole and screw-in-hole tasks are one of the commonest types of assembly task. This suggests that where the environment can be controlled, such as in assembly operations where the handled components are man-made and deterministic in structure, robotic automation may be employed, at least in part. Assembly is one of the most labour-intensive tasks so its full or partial automation would represent a considerable ease on astronaut resources. In-orbit servicing is thus one space arena suited to this approach as this typically involves the manipulation of well-characterised, engineered objects. Interaction with the human teleoperator occurs only when high-level commanding is required at the task or sub-task level which has a characteristic timescale of seconds. This approach then would be suitable for in-orbit robotics up to geostationary orbits, and perhaps for limited geological exploration of the Moon such as was achieved by the Russian Lunakhod rovers.

5. In-Orbit Servicing by Man and Machine

To-date, in-orbit servicing tasks have been performed by astronauts, aided by the Shuttle Remote Manipulator System for retrieval and deployment of spacecraft. Astronauts are exposed to hazards every time they undertake extravehicular activity (EVA), particularly for strenuous and arduous work like in-orbit servicing. Hazards are imposed by a hostile, radiation-filled vacuum with the ever-present possibility of EMU (extravehicular mobility unit) suit rupture due to snagging, micrometeoroid or debris impacts, exposure to cryogens or corrosive fuels. EVA tasks are restricted by access, risk and complexity. All tasks require foot restraints and handholds, and loose objects such as tools require tethers to attachment points. The need for restraints limits the flexibility of the human manual worker in space. Astronauts need to be thoroughly trained in such activities to ensure that they do not exceed physiological limits which can be achieved rapidly due to the strenuous nature of EVA. EVAs are medically limited to 7 hours (of which 6 hours is useful work) and three EVA shifts per week. Human performance is limited by strength, vigilance, fatigue and reaction speed.

However, the chief limitation is on astronauts' limited space accessibility to only low inclination Low Earth Orbits (LEO). The Shuttle no longer operates from the Eastern Test Range at Vandenberg Air Force Base which gave it polar orbit access. This leaves no EVA capability in high inclination polar orbits the orbit of choice for Earth Observation satellites. Moreover there is currently no access to the geosynchronous orbits where most communication satellites reside. NASA has repeatedly given its justification for manned spaceflight on the basis of the astronauts' in-orbit servicing role, the Solar Maximum Repair Mission (STS 41-C in 1984) and the first Hubble Space Telescope repair mission (STS-61 in 1993) being spectacular examples. Given astronauts' limited coverage of orbits, the expense of their deployment and the exposure of their lives to hazards, this use of astronauts as manual "cable repairmen" does not constitute an effective argument in favour of manned spaceflight, nor is it an effective use of highly trained astronauts. These tasks are best done robotically if possible. The Solar Maximum Repair Mission is instructive – the two main tasks were the replacement of the spacecraft's Attitude Control System Module which was a designed-for-servicing task, and the Main Electronics Box exchange which was not designed for servicing (Davis, 1987, Adams et al., 1987). The first procedure took 35 minutes by EVA, 15 minutes when executed by laboratory teleoperation, and 40 minutes when executed by laboratory automation. The second task took 2 hours of EVA, 3 hours by laboratory dual arm teleoperation, and has yet to be achieved by laboratory automation, although automated task sequence planning has been simulated (Sanderson et al., 1988). The chief difficulty in this last task was in the handling of flexible extended objects such as thermal blankets, though with better tactile sensors and more sophisticated force control algorithms available today, this now looks feasible (Zheng and Paul, 1985; Hogan et al., 1985; Kopf, 1989).

NASA has two potential robotic in-orbit servicing programmes that it is exploring. Ranger is a developmental NASA programme to provide in-orbit robotic capability for Earth orbiting satellite servicing operations, to be controlled telerobotically from the ground (Parrish and Akin, 1996). It is similar in scope to the UK ATLAS (Advanced TeLerobotic Actuation System) design (Ellery, 1996, 2000). Robonaut is a developmental NASA programme with the express purpose of reproducing an EVA-suited astronaut capability using EVA astronaut toolkits in a teleoperated mode of control from within the ISS (E. Aldridge, 1999, private communication). Astronaut toolkits are basically power tools which apply torques to bolts, etc. to minimise wrist movements. Both programmes are in their early stages of development but if they do become operational the lessons learned and technological capabilities gained would be vast, opening up a role for robots as construction workers, be they teleoperated, telerobotic or automated. In addition, such robotic in-orbit services would comprise an extremely useful part of the near-Earth space infrastructure in extending the lifetime of existing space assets and minimisation of infant mortalities through battery module replacement, refuelling, cryogen replenishment and orbital replacement unit exchange. This would relieve astronauts to perform tasks that cannot be performed robotically such as complex tasks requiring high dexterity and/or high degrees of human decision-making such as scientific investigation and experimentation onboard the Shuttle and/or ISS. Such in-orbit services would be also suitable for commercialisation by the private sector (Ellery, 1996, 2000).

6. When is Artificial Intelligence Going to Happen?

Robotic autonomy through artificial intelligence (AI) is highly desirable as it reduces ground station requirements and costs which are a significant fraction of the operating costs of space missions. Further, if such techniques were sufficiently advanced to implement partial, or even complete, human level intelligence, then instead of exposing astronauts to the dangers and rigours of spaceflight, we could send out machines to do human level tasks and employ human level decision-making remotely in space and on planetary surfaces. This is the ultimate goal of AI. Such problem-solving capabilities are typified by the planning process which acts at the highest hierarchical level of robotic control. Planning is essentially thinking ahead to solve problems in simulation before acting them out. This ability to plan implies the ability to predict the consequences of one's actions in the future. Further, it means that feedback to the human supervisor does not have to occur

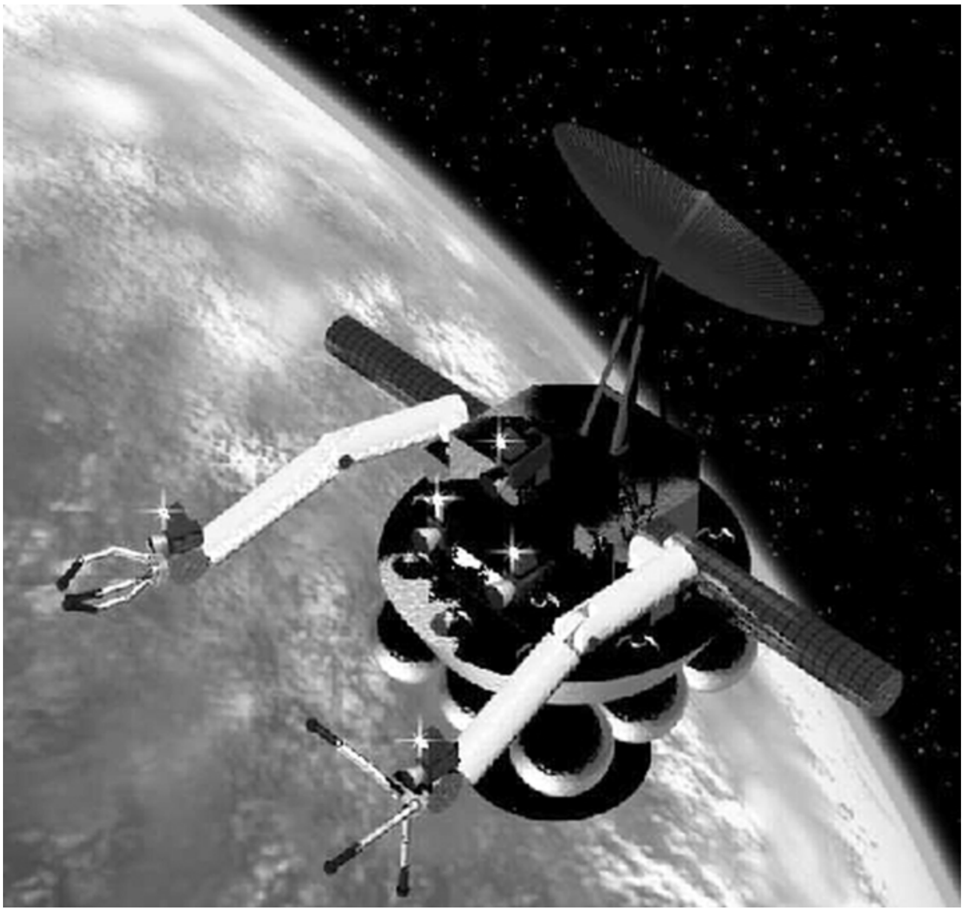

Figure 2. ATLAS in-orbit servicing robotic freeflyer.

very frequently as it does with teleoperated machines. To that end, robotic systems provide the most demanding test of AI techniques since it requires the robotic system to interact with the real world – sensory inputs and motor outputs are not analogous to standard read/write programming commands. There is a strong coupling of the robotic agent with its environment and it is this aspect of AI which is often ignored.

Early expectations in AI of "mechanised thought" spurred by advances in games playing and mathematical reasoning have yet to be realised in 50 years of research. All traditional AI methods are rooted in logic programming in which all intelligent behaviour is considered to be the result of symbol manipulation according to logical rules – planning is thus seen as an algorithmic problem-solving process which generates a sequence of actions to achieve certain objectives (goals). Planning may be regarded as a state-space search for appropriate actions to perform when applied to the current state of the world (as encoded in the world model) to change that state to another state that is closer to the goal state. The states of the world are modelled as logic statements in Horn clause logic such as Prolog which is computationally equivalent to the universal Turing machine. The current state of the world is analysed (the problem) and compared with the desired final state of the world (the solution) to find a procedure to reduce the difference (Charniak and McDermott, 1985). This effectively breaks an overall goal into smaller tasks to be achieved in order to realise the goal – this is sub-goaling in which the sequence of subgoals comprise the plan. This means-end analysis is applied recursively until the goal has been achieved – this was the basis of STRIPS (Stanford Research Institute Problem Solver) which is the fundamental basis of almost all automated planning systems (Fikes et al., 1972; Simon, 1991). Indeed, STRIPS was the basis of the Voyager spacecraft mission sequencing planner DEVISER (Vere, 1983). As the world model increases in size, the search space for a procedural path grows exponentially – it is an NP hard problem typified by the “Travelling Salesman Problem”. This is the computational explosion. The search space however can be “pruned” by implementing task-dependent production rules of the form “*if* this precondition is satisfied, *then* this action is effected” – these rules encode task-dependent knowledge that model human decision-making. This is the basis of the expert system in which large numbers of such rules (typically up to several thousand) encoding specialised human knowledge are constructed (Hayes-Roth, 1984, 1985). The expert system effectively comprises a world model. This approach has had limited applicability, though successes have been obtained in highly specialised and limited fields of application: PROSPECTOR for mineral prospecting; ACE for electronic component troubleshooting; STARPLAN for the diagnosis and correction of satellite malfunctions; Remote Agent planner for autonomous navigation onboard Deep Space 1. As expert systems encode semantic knowledge of limited applicability, outside of their limited domains of discourse they are useless - this is known as brittleness. These AI methods lend themselves to highly specialised applications involving well-structured environments such as in-orbit servicing where they could provide expert advice on planning assembly-disassembly operations. Their utility for planetary rovers would be limited without very extensive knowledge and controllability of the planetary environment in which the rover is operating. But this is precisely why planetary exploration is undertaken – to gain that knowledge.

Traditional AI systems utilise deductive inferencing and cannot reason inductively. Inductive reasoning involves generalising concepts from positive and negative examples. The ability to generalise through extracting important common features from a set of observations is a fundamental component of learning. This kind of capability is offered by artificial neural networks (ANN), but they are limited in scalability, particularly due to the long training times invoked by the backpropagation learning algorithm in large ANNs. Furthermore, there is no known way to algorithmically define the size and architecture of ANNs suited to a given problem. AI systems are capable of generalising in a highly limited fashion by converting constants in production rules into variables. This provides a highly limited form of analogical reasoning. A major flaw is that these limited capabilities are dependent on the human programmer rather than by direct interaction with the world (brain-in-a-vat criticism). Any interaction with the real world imposes the requirement that the environment is well behaved. The chessboard is an example of such a well-behaved environment. The limitations of logic programming have spurned a veritable industry of different methods for overcoming the limitations of logic manipulation – non-monotonic, default, modal, temporal and fuzzy logics to name a few – but none have solved the basic problems.

A brief but illustrative comparison is in order to compare human capabilities and machine limitations. Human-level intelligent capability implies the need for a vast repository of encyclopaedic knowledge about the world which humans acquire through learning (experience). This is why astronauts undergo lengthy periods of training. The amount of time required to program a robot with the equivalent of human intelligence would be very large: the human brain has an information capacity of 10^{14} bits in its 10^{11}neurons each with 10^3 synaptic connections. To program a machine with 10^{14} bits at one line of code per hour (the typical rate of software production from conception to installation) with an average line of code some 500 bits long would require 100 million man-years. Furthermore, on average there is around one bug per 100 statements in a program even after debugging, and this average increases with program size. In comparison, a human being typically learns 10^{14} bits through information input to the eyes and other senses at 250 Mbps over a 20 y period based on the brain's pattern recognition capabilities. Humans communicate through language, a symbol manipulation system that is a superset of logic unrestricted by the limitations of logical symbol manipulation. This has inspired the MIT Cog project to design and build a socially-aware, learning humanoid robot, but the research is still in its early stages.

7. Mars Exploration by Robotic Rovers

Planetary surfaces are highly unstructured environments and susceptible to uncertainties which must be accommodated in real-time. Longer-range missions to planets such as Mars incur time delays in signal propagation between Earth and the planetary surface – the round-trip communications time delay for Mars is 30 minutes on average. This precludes any form of teleoperative or telepresence means of direct human control. The Mars Pathfinder lander had only two five-minute communications windows to Earth per day. This was due to a lack of a Mars orbiter relay spacecraft which also restricted the Sojourner rover's mobility to within line-of-sight communications with the Pathfinder lander. The lander thus implemented a time-consuming store-and-forward communications strategy. Mobility as exemplified in planetary rovers introduces uncertainty and ensures that the environment changes as it moves, which must be constantly monitored. The

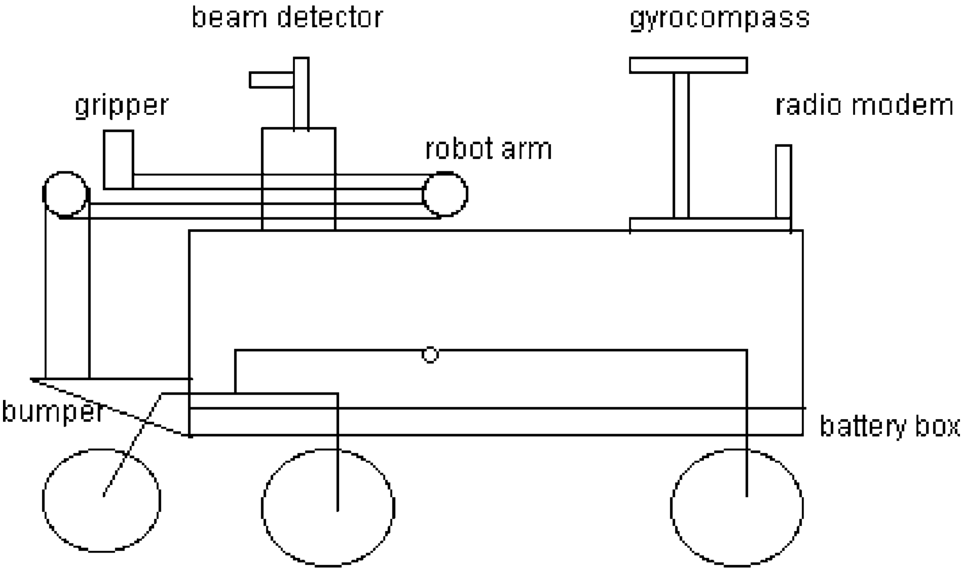

Figure 3. Rocky III/IV Mars Rover Testbed.

JPL Rocky series provides the model for planetary rover developments. They are limited to 10–100 kB behaviour control programs stored in ROM requiring up to 1 MIPS processing speeds. Rocky IV, a six-wheeled 15 kg microrover was the model for the Mars Sojourner rover.

Planetary rovers such as the Sojourner rover must be able to traverse unknown ground terrain with its attendant problems of steep slopes, crevasses, large rocks and loose soil. These conditions cannot be predicted in advance as orbital surveyor spacecraft are limited to images of the ground with ∼10 to 100 m resolutions. Waypoint navigation involves planning local routes autonomously with global routing being planned by the Earth operators. The human operator selects an approximate corridor to avoid large obstacles obtained from the lander mast-mounted panoramic stereo cameras while the rover implements reactive “behaviour” control based on its proximity sensor data locally. This was implemented on Sojourner and it was programmed to stop periodically to scan the route for obstacles and then continue. It was also programmed to stop and wait if it encountered tilting or wheel slippage – this caused extensive problems during its traverse of the “rock garden” indicating the inefficiency of this approach. In total, the vehicle traversed a total of 106 m within a 10 m radius of the lander over a three-month period through 114 commanded movements – an average of around 1 m per day. The primary role for a rover is as mobile science platforms, and Sojourner’s mission was highly successful. Sojourner implemented an alpha-proton-X ray backscattering absorption spectrometer (APXS) to analyse Martian soil and rock. The alpha-particle source was integrated into the sensor head which was deployed robotically against the sample due to the limited transmission distances of alpha particles. The typical integration time was 600 minutes per sample imposing a time-consuming static operation on the Sojourner's mobility.

The next mission to Mars is the Beagle II lander scheduled for 2003. Beagle II will carry a microscope, a gas chromatograph/mass spectrometer (GCMS similar to that used on Viking) to measure carbon isotope ratios, Mossbauer and X ray spectrometers integrated into the Beagle II arm. Its primary mission is to search for fossil evidence of life on Mars. The arm acts as an instrument placement device but will also implement core drilling of Martian rock for analysis of evolved gases in the lander's GCMS instrument following stepped pyrolysis. The GCMS instrument is the main instrument designed to detect any signs of organic material from the Martian regolith. Core drilling introduces the problem of vibration suppression – the most modern geological drills operate like jackhammers rather than rotary operation to eliminate drill-walk (Bar-Cohen et al., 2000). Beagle II will also carry a robotic "mole" to take sub-surface samples from up to 1 m depths for analysis in the GCMS instrument on the lander.

Future rover missions to Mars are presently dominated by the US effort involving the delivery of two identical 130 kg Athena-class rovers to two different sites on the surface of Mars also in 2003. Their mission profiles will resemble that of Mars Pathfinder in using airbags to impact the surface. Each rover will be able to traverse up to 100 m/day over an operational lifetime of 90 days. Their primary mission as "robotic field geologists" will be mineralogical characterisation, rather than the search for fossil life. They will carry an APXS like Pathfinder, a Mossbauer spectrometer like Beagle II and a thermal infrared spectrometer in a similar arrangement of a robotically deployable sensor head for long-duration spectroscopic integration times. The Athenas will also each include a Rock Abrasion Tool (RAT) for exposing fresh rock surfaces and a microscopic camera.

8. Planetary Rover Intelligence Quotient

The Good Old Fashioned Artificial Intelligence (GOFAI) approach to robotic autonomy based on logical theorem proving has met with limited success. Low level skills such as pattern recognition, learning and robust behaviours have not been adequate – a typical mobile rover implementing AI planning can take several minutes to act on an image frame. The Animat approach based on simple reactive behaviours offers a much more robust approach (Brooks, 1990; Wilson, 1990; Beer et al., 1990). This situated-robotics approach involves instigating simple reactive procedures to simple sensory signals. Examples of such reactive behaviours include generalised wandering behaviour, obstacle-avoidance behaviour and target-following behaviour each of which is selected on the basis of environmental cues. Brooks (1986, 1989) developed a "subsumption" architecture to implement each behavioural module which operated in parallel and implement the behaviours independently of each other. Each behavioural module was connected directly to sensors and actuators doing away with world models altogether. Such behaviour control is adequate for navigation with reactive obstacle avoidance in rovers, but has not been demonstrated for more complex tasks. Other methods are being explored such as genetic algorithms, genetic programs, neural networks and variations thereof, but it is not clear how these methods scale up to generate complex task-oriented behaviours (de Garis, 1991; Koza, 1993). By their very nature, they are not well-understood and imply a degree of uncontrollability. These situated robotics methods are based on the notion of robust survivability and less on the performance of more complex, goal-oriented tasks. A truly autonomous robot capable of performing useful tasks in the absence of direct human control requires both goal-oriented planning capability with a robust reactive capability – this robotic implementation of intelligence (the "sybaritic" approach to AI) is where the true test of AI must be (Ellery, 2000). Research into integrating different approaches to AI to give the robot a complement of capabilities is a very active area of research at present but is proceeding very slowly, e.g., Mataric (1992).

There have been some successes in developing robotic autonomy. Nomad is a 0.5 tonne autonomous rover capable of traversing 10–20 km of rough terrain over a month (Whittaker et al., 1997). It is currently in service in the Elephant Moraine rock-field in the Antarctic, autonomously searching for meteorites in the ice. It is worth noting that it is powered by a petrol-driven electric generator which supplies 2 kW of power which would require fundamental redesign for service on Mars. Furthermore, its vision processing is performed at a ground station, precluding its use on remote planetary environments. Not all rover missions have had such success. The 800 kg Dante rover was an eight-legged robotic rover designed to descend into active volcanic craters in the Antarctic, teleoperated through a tether cable. It was designed to minimise dangers to field geologists, eight of whom had been killed in 1993. Its maiden mission was characterised by slow performance due to limitations in the control system until its fibre optic tether snagged and broke. State-of-the-art robotics technologies such as artificial intelligence are insufficient to replace human intelligence in the field, particularly on unstructured planetary environments. Robots must interact with the real physical world which is imprecisely modelled, full of uncertainty and subject to change. The only way to implement autonomy on spacecraft is to minimise the environmental complexity within which it operates, but this is not possible for all space missions such as planetary rover missions. For comparison, Deep Space 1 Remote Agent's environment was spacecraft navigation – a very well-understood environment. By the very nature of planetary exploration, we cannot predict in advance the incidence of all possible events that might occur and all possible failure modes. Planetary environments are only partially known and characterised, otherwise we would not explore them.

9. Humans Step in and Save the Day

The deployment and conduct of scientific experiments and sample collection are often complex and require real-time analysis by the human astronaut. Human dexterity backed up with human inductive decision-making is required to install and maintain complex scientific equipment. The maintenance of such equipment can be complex as maintenance is essentially a form of repair mitigation, reacting to sub-optimal performance prior to failure. When complex equipment does fail, it often requires human intelligence and dexterity to repair, particularly if it has moving parts, sometimes requiring complex workarounds that would not be possible robotically.

Human beings are ideally suited to the conduct of detailed, field exploration which require flexibility, skill and judgement beyond that of machine capabilities (Crawford, 1998; Spudis, 1999). Exploration is an iterative process whereby as exploration proceeds, data is collected and quick-look analysed to provide interpretation which further constrains additional exploration. This process cannot be conducted robotically by rovers – human field geologists will be required to select sites for study and on the basis of their findings decide further exploration strategies. If such exploration were conducted robotically, each time a discovery was made it would require a further robot mission with more tightly constrained sensors to capitalise on the results of the previous robotic mission. This process would likely require at least many decades to produce the desired results as launch opportunities arise only once every two years (Zorpette, 2000).

We have on Earth only 16 SNC meteorites from Mars which are undergoing extensive scientific study at present. This invalidates the concept of a robotic Mars sample return mission, as this would add little to the inventory of Mars material on Earth at great cost. However, a human mission could recover much greater amounts (from the lunar example, around 300 times) making it far more cost-effective.

If the division of labour between man and machine is designed thoughtfully, then the scientific returns from any space exploration mission can be maximised. To that end, further Mars lander and rover missions are required to support a series of manned landings on Mars (Zubrin and Weaver, 1993, 1995; Zubrin et al., 1991). However, it is important to remember the fate of the Space Exploration Initiative (SEI) supported by President George Bush. Its price tag of $30 B/y over 15 years to give a total expenditure of $500B effectively killed it. Furthermore, it was marketed as an adventure not as a scientific endeavour. The situation today however is very different, especially with the discovery of the ALH84001 SNC meteorite and the proposal that it offers evidence for life on Mars during its early history (McKay et al., 1996). Although the significance of the find is still being debated, the recent discoveries of possible ancient lake-beds and perhaps more recent water flows on the surface of Mars by the Mars Global Surveyor spacecraft have lent support to the notion of life on Mars during its early history, and perhaps more recently.

The search for life on Mars, be it fossil or extant, must be the most important goal for all space exploration. Indeed, it is one of the most important questions for science as a whole. The evidence for such life on Mars is likely to be sparse, microscopic, hidden and difficult to access – the proverbial needle in a haystack. Robotic rover missions, limited as they are, will not be up to the task – they lack the necessary complex capabilities for core site exploration, deep coring and core manipulation. Their job is to survey and gain panoramic knowledge of the Martian environment. A thorough survey of Mars could be undertaken by a buckshot spread of perhaps 100 microrovers of 5–10 kg in mass to increase the area coverage over the surface with a high degree of redundancy, but the mechanism of delivery to the surface is not clear, and the expense is likely to be high. Planetary rovers at present are only capable of obtaining soil samples close to the surface – at the most optimistic, a couple of metres into the regolith. Humans in the field are the only option in providing high mobility, wide area coverage, capability to dig to great depths and most importantly, high adaptability and flexible decision-making. Only humans can make informed and expert judgements based on their past professional experiences and build "intuitive" analogies based on those experiences. Analogy involves the use of prototypical situations built from experience to allow recognition of the present situation in terms of past situations. Only humans have this capability. Human experience is built up in episodic memory while AI methods are almost entirely procedural and prescriptive.

Manned missions to Mars today have a very clear and unequivocal scientific goal, unlike the SEI Mars vision. With this in mind, it is essential that a human mission to Mars is the only solution to this goal, supported by an adequate degree of robotics technology. In support of human missions, robotics technologies will provide the means for serial precursor robotic rover survey missions for surface characterisation, the automated manufacture of propellant on the surface to minimise costs, automated and/or teleoperated manual robots for habitat deployment on the surface, multiple teleoperated/manned rover surveys around the landing site and beyond, and automated scientific laboratories to ease the workload of the astronaut-scientists. These technologies could be developed and field-tested through a series of unmanned and/or manned lunar missions which impose lesser constraints on remoteness than Mars. Human missions to Mars are more likely to succeed on the basis that more attention is paid to quality assurance, robustness and safety than in unmanned missions. The track record of robotic missions to Mars is highly variable with spectacular successes interspersed with catastrophes. Although the recent loss of the Mars Polar Lander and Mars Climate Orbiter have been attributed to insufficient funding due to the cheaper part of the "faster, cheaper, better" philosophy and a lack of effective project management, this cannot be attributed to the loss of Mars Observer and Mars-96. A human mission to Mars cannot be subjected to the same constraints due to potential for catastrophic loss of life.

The discovery of fossil or extant life on Mars, if it does indeed exist, will be only the first step in the process. Although important in itself philosophically, this information will be scientifically useless unless it is followed up by a comprehensive programme of scientific study. Repeatability is a cornerstone of scientific rigour. This means attempting to unveil its distribution, its variety, its habitats, its origin, its viability, etc. And this cannot be accomplished by robotic probes alone.

10. Conclusion

I have argued that there are fundamental problems in the approach to space exploration undertaken by the various space agencies with regard to the division of labour between man and machine. These issues should be more carefully designed with a symbiotic approach in mind. I argue that space robotics is an enabling technology to complement human activities, but it has its limitations. State-of-the-art artificial intelligence is insufficient to replace astronauts in the field on planetary surfaces such as Mars. This assessment suggests that given the importance of the search for life on Mars, humans will have to go to Mars and more, maintain a concerted and sustained study of the Martian environment.

Human beings have only recently discovered life on our own planet in extreme environments, the Archaea. The classification tree of life on Earth that has been entrenched in biology textbooks since Linnaeus is being rewritten on the basis of these recent discoveries. It will only be through extensive manned scientific study of Mars that we will be able to rewrite the tree of life in the Universe as a whole – this will be one of the most important quests in modern science with vast implications across a large number of scientific disciplines.

References

Adams, R. et al.: 1987, ‘Remote Repair Demonstration of Solar Maximum Main Electronics Box’, in *Proceedings of the First European In-Orbit Operations Technology Symposium* (ESA SP-272), pp. 227–323.

Bar-Cohen Y. et al.: 2000, ‘Ultrasonic/Sonic Drilling/Coring (USDC) for *in-situ* Planetary Applications’, *SPIE Smart Structures 2000*, Paper No. 3992-101.

Beer R. et al.: 1990, ‘Biological Perspective on Autonomous Agent Design’, *Robot. Auton. Syst.* **6**, 169–186.

Bejczy, A.: 1980, ‘Sensors, Controls and Man–Machine Interface for Advanced Teleoperation’, *Science* **208**, 1327–1335.

Brooks, R.: 1986, ‘Robust Layered Control System for a Mobile Robot’. *IEEE Trans. Robot. Autom.* **2**, 14–23.

Brooks. R.: 1989, ‘Robot that Walks: Emergent Behaviours from a Carefully Evolved Network’, *Neural Comput.* **1**, 253–262.

Brooks, R.: 1990, ‘Intelligence without Representation’, *Artif. Intell.* **47**, 159–189.

Charniak, E. and McDermott, D.: 1985, *Introduction to Artificial Intelligence*, Addison-Wesley Publishers, Reading, MA.

Cohen, A. and Erickson, J.: 1985, ‘Future Uses of Machine Intelligence and Robotics for the Space Station and its Implications for the US Economy’, *IEEE Trans. Robot. Autom.* **1**, 117–123.

Crawford, I. A.: 1998, ‘The Scientific Case for Human Spaceflight’, *Astron. Geophys.* **39**, 6.14–6.18.

Davis, R.: 1987, ‘In-Orbit and Laboratory Exchange of ORUs Designed/Not Designed for Servicing’, in *Proceedings of the First European In-Orbit Operations Technology Symposium* (ESA SP-272), pp. 123–126.

De Garis, H.: 1991, ‘Genetic Programming: Building Artificial Nervous Systems with Genetically Programmed Neural Network Modules’, in B. Soucek (ed.), *Neural and Intelligent Systems Integration*, John Wiley Publishers, New York, NY, pp. 207–234.

Ellery, A.: 1996, ‘Systems Design and Control of a Dual-Manipulator Freeflying Robotic Spacecraft for In-Orbit Satellite Servicing’, Ph.D. Thesis, Cranfield University, UK.

Ellery, A.: 2000, *An Introduction to Space Robotics*, Praxis-Springer Publishers, UK.

Fikes, R., Hart, P., and Nillson. N.: 1972, ‘Learning and Executing Generalised Robot Plans’, *Artif. Intell.* **3**, 251–288.

Hayes-Roth, F.: 1984, ‘Knowledge-based Expert Systems: A Tutorial’, *IEEE Comp.* (Jan), 9–24.

Hayes-Roth, F.: 1985, ‘Rule-Based Systems’, *Comm. Assoc. Comput. Mach.* **37**, 27–39.

Hogan, J. et al.: 1985, ‘Impedance Control: An Approach to Manipulation – Part I–III’, *ASME J. Dynam. Syst. Measur. Contr.* **107**, 1–24.

Kopf, C.: 1989, ‘Dynamic Two Arm Hybrid Position/Force Control’, *Robot. Autonom. Syst.* **3**, 369–376.

Koza, J.: 1993, ‘Evolution of Subsumption Architecture Using Genetic Programming’, in *Proceedings of the First European Conference on ALife*, pp. 110–119.

Kuperstein, M. et al.: 1987, ‘Adaptive Visual-Motor Coordination in Multi-Jointed Robots Using Parallel Architectures’, in *Proceedings of the IEEE International Conference on Robotics & Automation*, pp. 1595–1602.

Mataric, M.: 1992, ‘Integration of Representation into Goal-Driven Behaviour-Based Robots’, *IEEE Trans. Robot. Autom.* **8**, 304–312.

McKay, D. et al.: 1996, ‘Search for Past Life on Mars: Possible Relic Biogenic Activity in Martian Meteorite ALH84001’, *Science* **273**, 924–930.

Parrish, J. and Akin, D.: 1996, ‘Ranger Telerobotic Flight Experiment: Missions, Technologies and Pragmatics’, in *Robotics for Challenging Environments: Proceedings of the RCEII Conference*, pp. 136–142.

Rouse, W. and Cody: 1987, ‘On the Design of Man-Machine Systems: Principles, Practices and Prospects’, in *Proceedings of the IFAC 10th Triennial World Congress*, Munich, pp. 281–288.

Sanderson, A., Peshkin, M., and Homem-de-Mollo, L.: 1988, ‘Task Planning for Robotic Manipulation in Space Applications’, *IEEE Trans. Aero. Elect. Syst.* **24**, 619–628.

Schmidt, G. and Hawes, M.: 1999, ‘Robots vs Humans in Space: BOTH Will Be Required’, *Ad Astra* (http://www.astrobiology.com/adastra/robots.vs.humans.html)

Simon, H.: 1991, ‘AI: Where It Has Been and Where Is It Going?’, *IEEE Trans. Knowl. Data Engin.* **3**, 128–136.

Spudis, P.: 1999, ‘Robots v Humans in Space’, *Sci. American Presents: The Future of Space Exploration* **10**, 25–31.

Vere, S.: 1983, ‘Planning in Time: Windows and Durations for Activities and Goals’, *IEEE Trans. Pattern Anal. Mach. Intell.* **5**, 246–260.

Whittaker, W. et al.: 1997, ‘Atacama Desert Trek: A Planetary Analog Field Experiment’, in *Proceedings of the International Symposium AI, Robotics & Automation for Space*, Tokyo, Japan.

Wilson, S.: 1990, ‘The Animat Path to AI’, in *Proceedings of the First International Conference on Simulation of Adaptive Behaviour (From Animals to Animats)*, Paris, pp. 15–21.

Zheng, Y. and Paul, R.: 1985, ‘Hybrid Control of Robot Manipulators’, in *Proceedings of the IEEE International Conference on Robotics and Automation*, pp. 602–608.

Zorpette, G.: 2000, ‘Why Go to Mars?’, *Sci. Amer.* (March).

Zubrin, R. and Weaver, D.: 1993, ‘Practical Methods for Near-Term Piloted Mars Missions’, *AIAA 91-2089.*

Zubrin, R. and Weaver, D.: 1995, ‘Practical Methods for Near-Term Human Exploration of Mars’, *J. Brit. Interplanet. Soc.* **48**, 287–300.

Zubrin, R. et al.: 1991, ‘Report on the Construction and Operation of a Mars *in-situ* Propellant Production Unit’, *J. Brit. Interplanet. Soc.* **48**, 327–336.