Challenges and Opportunities for Autonomous Systems in Space

Chad R. Frost

Supervisor, Autonomous Systems and Robotics

Intelligent Systems Division

NASA Ames Research Center

Moffett Field, California

U.S. Frontiers of Engineering Symposium
September 22, 2010

Abstract

Despite notable successes establishing the viability of spacecraft autonomy, widespread adoption has been slow in coming. The missions that have paved the way have also helped clarify the key challenges to broader use of autonomous systems. There are many opportunities for autonomy to enable or enhance space missions if these challenges can be overcome.

1 Introduction

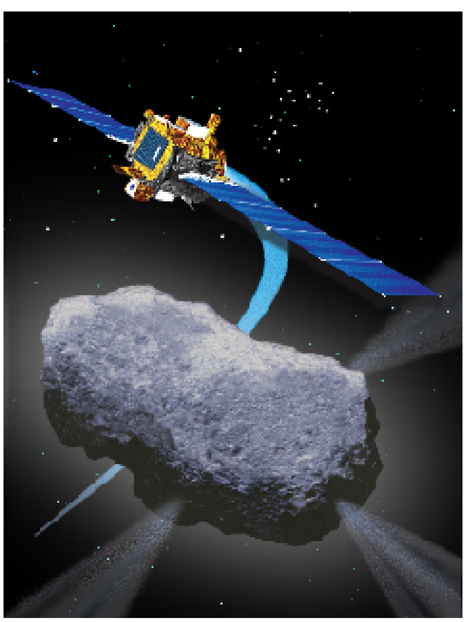

Figure 1: Deep Space 1 flew the Remote Agent Experiment, demonstrating full spacecraft autonomy for the first time. [1]

1 All figures courtesy NASA except as noted.

With the launch of Deep Space 1 in 1998, the autonomous systems community celebrated a milestone – our first flight experiment demonstrating the feasibility of a fully autonomous spacecraft. We anticipated that the kind of advanced autonomy demonstrated on Deep Space 1 would soon be pervasive, helping to enable science missions, make spacecraft more resilient, and reduce operations costs.

However, the pace at which autonomy has found its way onto space missions has been relatively slow. Where it has been used, either operationally or as an experiment or demonstration, autonomy has succeeded; outstanding work by the community has continued to advance the technologies. [Chien et al., 2005, Castaño et al., 2006, Estlin et al., 2008, Fong et al., 2008, Knight, 2008] There are numerous drivers for putting autonomy on board spacecraft – including maintenance of the spacecraft despite failures or damage, extension of the science team through "virtual presence", and cost-effective operation over long periods of time. Why, then, has autonomy not been more broadly adopted?

Definition of Autonomy

It's useful to clarify the difference between *autonomy* and *automation*. Many definitions are possible (an example being [Doyle, 2002]), but here we'll focus on one of the common needs of systems outside our direct, hands-on control: the need to make choices.

An "autonomous system" is defined as a system that resolves choices on its own. The goals the system is trying to accomplish are provided by another entity; thus, the system is independent from the entity on whose behalf the goals are being achieved. The decision-making processes may in fact be simple, but the choices are made locally.

In contrast, an "automated system" follows a script, albeit a potentially quite sophisticated one; if it encounters an unplanned-for situation, it stops and waits for human help, e.g. "phones home". The choices have either been made already, and encoded in some way, or will be made externally to the system.

An "intelligent" autonomous system chooses using more sophisticated mechanisms – often resembling those used by humans, but regardless of implementation details, capable of more "creative" solutions to ambiguous problems. Ultimately, the intelligence of an autonomous system is judged by the quality of the choices it makes.

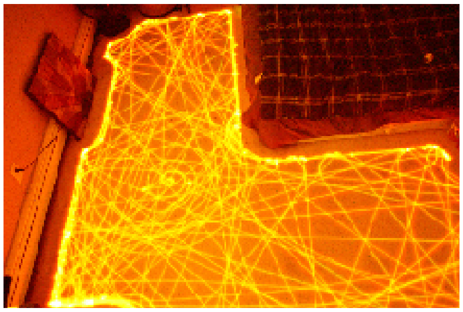

Figure 2: Long-exposure image of Roomba's path while navigating a room. Photo © Paul Chavady, used with permission.

A domestic example that we've probably all seen is iRobot's Roomba™ line of robotic vacuum cleaners – they illustrate just how prosaic autonomous systems have become. Navigating a houseful of obstacles while ensuring the carpet gets cleaned is a pretty challenging task for a consumer product. Using the definitions above, we note:

- the Roomba user provides goals (vacuum the floor, but don't vacuum here, vacuum at this time of day, etc.)
- the Roomba has choices to make (how to identify the room geometry, avoid obstacles, when to recharge its battery, etc).
- the Roomba also has some automated behavior and encounters choices it can't resolve on its own (it gets stuck, it can't clean its own brushes)

The Roomba may be an example of marginal autonomy – there are numerous situations it can't deal with by itself, and it's certainly not "intelligent" – but it does have basic on-board diagnostic capability ("clean my brushes!") and has a strategy (as seen in Figure 2) for vacuuming a room whose size and layout of which it initially is ignorant. Roomba™ demonstrates the potential for autonomous systems to become common in our daily experience.

2 Why is autonomy needed by space missions?

So much has been written on this topic that we can barely begin to scratch the surface of a deep and rewarding discussion. However, we can examine a few of the recurring themes. The needs depend, of course, on the mission. Autonomous operation of the spacecraft, its subsystems, and the science instruments or payload, becomes increasingly important as the spacecraft is required to deal with phenomena that occur on timescales smaller than the communication latency between the spacecraft and Earth. But all spacecraft have in common the need for maintaining function "to ensure that hardware and software are performing within desired parameters, and finding the cause of faults when they occur."[Post and Rose]

"Virtual presence" is often cited as one of the most compelling needs for autonomy. In particular, scientific investigation, including data analysis and discovery, becomes more challenging the farther away the scientific instruments are from the locus of control. Doyle suggests "…a portion of the scientist's awareness will begin to move onboard, i.e., an observing and discovery presence. Knowledge for discriminating and determining what information is important would begin to migrate to the space platform." [Doyle, 2002] AI pioneer Marvin Minsky commented in 1980,

> "With a lunar telepresence vehicle making short traverses of one kilometer per day, we could have surveyed a substantial area of the lunar surface in the ten years that have slipped by since we landed there." [Minsky, 1980]

Adding three decades to Minsky's sentiment only serves to reinforce it. Traversing extraterrestrial sites and more broadly, conducting the ongoing "Dirty, Dull, Dangerous" work that it is not cost-effective, practical, or safe for humans to do, is a strong and recurring need with great potential for long-term benefit. As pointed out by Minsky, even if the pace of investigation or work is not what humans *in situ* might accomplish, the cumulative effort can be impressive. The Mars Exploration Rovers are a fine example of this, having jointly traversed more than 18 miles and collected hundreds of thousands of images in the six years since they landed.

Another flavor of "virtual presence" is the extension of the spacecraft operator's knowledge onto the spacecraft, enabling "greater onboard adaptability in responding to events, closed-loop control for small body rendezvous and landing missions, and operation of the multiple free-flying elements of space-based telescopes and interferometers." [Doyle, 2002] Several of these aspects have been demonstrated, such as the use of Livingstone 2 to provide on-board diagnostics on EO1, Hayden et al. [2004] and full autonomy including docking and servicing with Orbital Express. Ogilvie et al. [2008]

In addition to the robotic mission enablers above, there is the need to reduce the workload of astronauts in crewed missions. Astronauts' limited and valuable time should not be spent verifying parameters and doing spacecraft housekeeping, navigation, and other chores that could readily be accomplished by autonomous systems.

These various needs distill down to a few common underlying themes: Mitigating latency (e.g. distance from those interested in what the system is doing), improving efficiency (e.g. through reduction in cost and mass, and improved utilization of instrument and/or crew time), and managing complexity (in that systems are now so complex that it can be difficult or impossible for humans, whether on-board or on the ground, to diagnose and solve problems). Nothing is flown on a spacecraft without having "paid" for itself, and this holds true for the software to give it autonomy as well.

3 Success stories

Where then do we stand today, in terms of *deployed* autonomy? Specifically, what technology elements that comprise a complex, self-sufficient, intelligent system have flown in space? There have been several notable successes.

Deep Space 1 - Remote Agent Experiment

In 1996, NASA specified an autonomous mission scenario called the New Millennium Autonomy Architecture Prototype (NewMAAP). A Remote Agent architecture, integrating constraint-based planning and scheduling, robust multi-threaded execution, and model-based mode identification and reconfiguration, was developed to meet the NewMAAP requirements. [Pell et al., 1996, Muscettola et al., 1998]

As described by Muscettola et al,

> "The Remote Agent architecture has three distinctive features: First, it is largely programmable through a set of compositional, declarative models. We refer to this as model-based programming. Second, it performs significant amounts of on-board deduction and search at time resolutions varying from hours to hundreds of milliseconds. Third, the Remote Agent is designed to provide high-level closed-loop commanding."

The success of the NewMAAP demonstration resulted in the Remote Agent being selected as a technology experiment on the Deep Space 1 mission. In 1998, NASA launched Deep Space 1 (DS1, shown in Figure 1), with the goal of testing twelve cutting-edge technologies including the Remote Agent Experiment (RAX) – which became the first operational use of "artificial intelligence" in space. DS1 functioned completely autonomously for 29 hours, successfully operating the spacecraft and responding to both simulated and real faults.

Earth Observing 1 - Autonomous Sciencecraft Experiment

Earth Observing 1 (EO1) was launched in 2000, and demonstrated on-board diagnostics and autonomous acquisition and processing of science data – specifically, imagery of dynamic natural phenomena that evolve over relatively short timespans, such as volcanic eruptions, flooding, ice breakup, and changes in cloud cover.

The EO-1 Autonomous Sciencecraft Experiment (ASE) included autonomous on-board fault diagnosis and recovery (Livingstone 2), as well as considerable autonomy of the science instruments and downlink of the resulting imagery and data. Following demonstration, ASE was adopted for operational use and has been in operation since 2003. [Chien et al., 2005, 2006]

One of the signal successes of the EO1 mission was its independently capturing volcanic activity on Mt. Erebus. In 2004, ASE detected an anomalous heat signature, scheduled a new observation, and effectively detected the eruption by itself. Figure 3 shows a 2006 collection from EO1 of the same volcano.

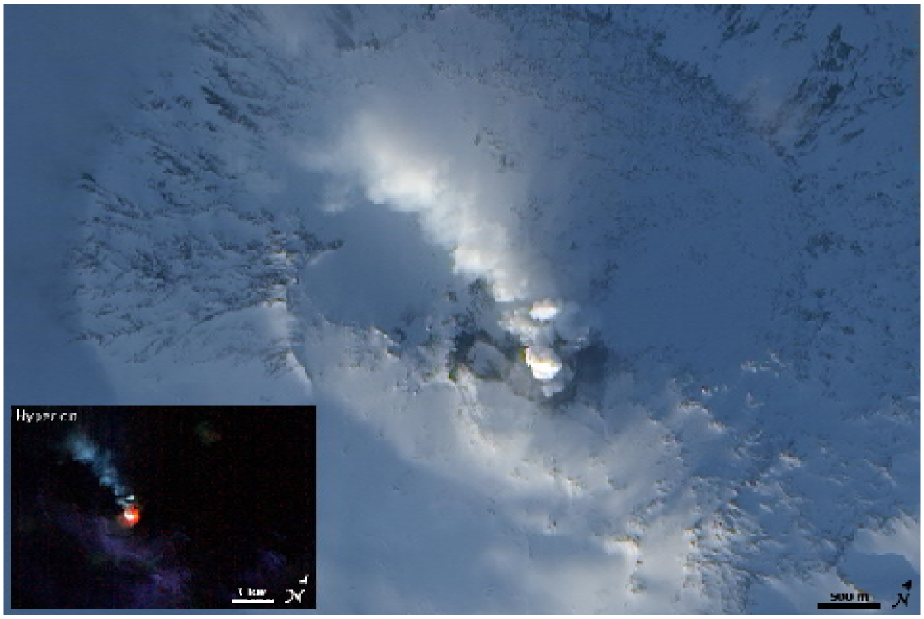

Figure 3: Images of volcanic Mt. Erebus, autonomously collected by EO1. NASA image created by Jesse Allen, using EO-1 ALI data provided courtesy of the NASA EO1 Team.

Orbital Express

In 2007, the Orbital Express mission launched two complimentary spacecraft, ASTRO and NextSat, with the goal of demonstrating a complete suite of technologies required to autonomously service satellites on-orbit. The mission demonstrated several levels of on-board autonomy, ranging from mostly ground-supervised up to fully autonomous capture and servicing, self-directed transfer of propellant, and automatic capture of another spacecraft using a robot arm. [Boeing Integrated Defense Systems, 2006] [Ogilvie et al., 2008] These successes demonstrated that autonomous servicing and other complex spacecraft operations can be conducted.

Mars Exploration Rovers

The Mars Exploration Rovers, Spirit and Opportunity, have had increasing levels of autonomy since their landing on Mars in 2004. An early enhancement provided autonomous routing around obstacles; another automated the process of calculating how far to reach out a rover's arm to touch a particular rock. In 2007, the rovers were updated to autonomously examine sets of sky images, determine which ones show interesting clouds or dust devils, and then to send only those images back to scientists on Earth. The most recent software enables Opportunity to make decisions about acquiring new observations – rocks imaged by the wide-angle navigation camera are selected by shape and color for detailed inspection with the narrow-field panoramic camera.[Estlin et al., 2008]

4 Challenges to Broader Use

With compelling needs for spacecraft autonomy, and feasibility demonstrated by the successful missions described in the previous section, what obstacles remain to autonomous systems' emergence as regular elements of spacecraft flight software? Two kinds of requirements for spacecraft autonomy can be articulated. The first are *functional* requirements, representing the attributes the software must objectively satisfy for it to be acceptable. A second group are *perceived* requirements: those that may not be grounded in real mission requirements, but have (possibly large) subjective weight. Both groups must be successfully addressed for autonomy to move into common use.

Functional Requirements

From our experience thus far, we have a good sense of what the overarching requirements for space mission autonomy actually are. Muscettola et al. [1998] offers a nicely distilled set of these requirements (evolved from [Pell et al., 1996]), which are quoted and expounded upon here.

"First, a spacecraft must carry out autonomous operations for long periods of time with no human intervention."

Otherwise, what's the point? Short-term autonomy may even be worse than none at all; if humans have to step in to what is nominally an autonomous process, they can spend a lot of time working to understand what the state of the spacecraft is, how it got there, and what needs to be done about it.

"Second, autonomous operations must guarantee success given tight deadlines and resource constraints."

By definition, the autonomous system can't just stop and wait indefinitely (either for human help, or to deliberate on a course of action) if an unplanned circumstance arises; it *must* act, and in an expedient fashion. Whether autonomy can truly guarantee success is debatable, but it has to at least provide the highest likelihood of success.

"Third, since spacecraft are expensive and are often designed for unique missions, spacecraft operations require high reliability."

Even in the case of a (relatively) lower-cost mission, with expressed acceptance of a higher level of risk, we tend to be quite risk-averse! Failure is perceived to be undesirable, embarrassing, "not an option" even for those missions that have in theory traded increased chance of failure for reduced cost or schedule. Therefore, autonomy has to be perceived as *reducing* risk, and ideally without significant increases to cost and schedule. Program managers, science PIs and spacecraft engineers want (and need) an answer to their oft-asked question, “How can we be sure that your software will work as advertised and avoid unintended behavior?”

“Fourth, spacecraft operation involves concurrent activity among a set of tightly coupled subsystems.”

This aspect requires that requirements and interfaces be thoroughly established relatively early in the design process; it pushes software efforts towards the front of the program, changing the cost profile of a mission.

Perceived Requirements

Opinions and perceptions, whether grounded or not, are significant challenges to flying autonomy on spacecraft. Some of the key requirements and the associated issues are:

Reliability.

As noted above, this is the most-asked question; autonomy software is perceived as adding a lot of potential risk to a mission.[Frank, 2008b] The question really being posed is “does the ability of autonomy to deal with otherwise loss-of-mission failures sufficiently offset the added potential for software problems?”

Complexity.

Bringing autonomy onto a spacecraft is perceived as adding a lot of complexity to what might be an otherwise fairly straightforward system. Certainly autonomy increases the complexity of a simple, bare-bones system. This becomes a more nuanced issue if we’re talking about *degrees* of autonomy, or autonomy versus automation. As spacecraft systems become more complex themselves, or we ask greater independence of them, does the software have to increase in sophistication to match? Does sophistication equal complexity?

Cost.

Adding many lines of code to support autonomy functions is perceived as driving up mission costs. As above, the potential for the autonomy to uniquely enable or save the mission is the primary way it “buys” its way onto a spacecraft. The opportunity for long-term savings in reducing operations costs should not be overlooked. However, in both cases, the benefits may be much more difficult to quantify than the costs; this is the aspect that can make the cost of deploying autonomy a highly subjective attribute.

“Sci-fi views of Autonomy.”

This comes down to our ability to manage expectations and educate those outside the intelligent systems community; autonomous systems are perceived as either overly-capable, borderline Turing machines, or latent HAL-9000s ready to run amok (or worse, to suffer from some more subtle and hard-to-diagnose form of mental illness.) Sorry, but we’re just not there yet!

5 How do we address the challenges?

Risk-reduction (perceived and otherwise) is the principal hurdle autonomous systems must overcome. Improvements could be achieved through:

Processes.
Perhaps the greatest potential contributor to regular use of autonomy on spacecraft is to ensure that rigorous processes are in place to (a) thoroughly verify and validate the software, (b) minimize the need to develop new software for every mission. Model-based software helps fulfill both key aspects (as well as others) – it allows rigorous validation of the component models, and re-use of *knowledge* even on one-off spacecraft development. The model-based software approach was used for the Remote Agent Experiment, [Williams and Nayak, 1996] and for the Livingstone 2 diagnostic engine used aboard EO1. [Hayden et al., 2004] However, processes do not emerge fully-formed; only through the experience of actually flying autonomy (encompassing both successes and failures) can we learn about our processes and the effectiveness of our methods. Thus, we need to engage in:

Demonstrations.
The next most effective way to retire risk is to increase the flight experience and heritage of the autonomy software components. This necessitates a methodical approach to including autonomy on numerous missions, as "ride-along" secondary or tertiary mission objectives initially but eventually to include missions dedicated to autonomy experiments. This is by no means a new idea. NASA launched DS1 and EO1 (and three other spacecraft) under the auspices of the New Millennium Program, which was initiated in 1995 with the objective of validating a slate of technologies for space applications. However, there is not presently a long-term strategy in place to continue development and validation of spacecraft autonomy.

Fundamental research. Autonomy has a long way to go to achieve all that we would have it do; considerable research is yet required to identify new, creative solutions that will address the many challenges. For example, planning and fault detection can both be difficult problems to solve; running on modern terrestrial computers they may take too long to solve, let alone when executing on much-slower flight-qualified hardware. Academic experts know how to create planners or FDIR software that can run fast enough, but this expertise needs to be transformed into engineering discipline. This has been accomplished successfully in related domains, as for example linear programming techniques that have been packaged to allow Operations Research students to learn how to build LP models using standard industrial tools.

Education.
Simultaneously with the above efforts, we need to reach out and actively educate Principal Investigators, project managers and the science community about the advantages offered by autonomy, and what the associated costs and savings actually are.

6 Opportunities

There are many upcoming missions in which autonomy could play a major role.

Human spaceflight has, to date, been remarkably devoid of autonomy. However, as the technologies are validated in unmanned spacecraft and reach levels of maturity commensurate with other human-rated systems, there is great potential for autonomy to assist the crew with maintaining and operating the most complex of spacecraft over long periods of time. Life support, power, communications, and other systems require automation, but would benefit from autonomy. [Frank, 2008a]

Missions to the planets of our solar system, and their satellites, will increasingly demand the greatest possible productivity and scientific return on the large investment represented by the development and launch of these sophisticated spacecraft. Particular destinations, such as the seas of Europa, will place great demands on autonomy to carry out independent exploration in environments where communication with Earth is difficult or impossible. Proposed missions to Near Earth Objects (NEOs) necessitate autonomous rendezvous and proximity operations, and possibly contact with or sample retrieval from the object.

A variety of Earth and space science instruments can potentially be made autonomous, in whole or in part, independently of whether the host spacecraft has any autonomy for operations. Autonomous drilling equipment, hyperspectral imagers, and rock samplers have all been developed and demonstrated terrestrially in Mars-analog environments. EO1 and the Mars Exploration Rovers were (and are) fine examples of science autonomy deployed to mitigate our inability to immediately respond to transient phenomena; there are certainly other possibilities.

These are by no means an exhaustive list – but representative of the broad spectrum of opportunities to put autonomy to good use in support of space missions. Numerous others, in space, on Earth, below our seas, could be enhanced or brought to reality through the careful application of autonomy.

7 Conclusions

Looking back, we've had some great successes, and the opportunities we have to look forward to are hugely exciting. But it has been more than a decade since the first autonomous systems in space, and operational autonomy is not yet a standard practice. What then would enable adoption of autonomy by more missions? We've argued in section 5 that we need:

1. The infrastructure (in particular the development, testing, and validation processes) and codebase in place, so that each new mission doesn't have to re-invent the wheel, and only has to bear marginal costs,

2. A track record ("flight heritage") establishing reliability,

3. Broader understanding of the benefits of autonomous systems.

How do we get there?

To build infrastructure and heritage, we need to invest. Earth-based rovers, submarines, aircraft and other "spacecraft analogs" can serve (and are frequently serving) as lower-cost, lower-risk validation platforms. Such developmental activities should lead to several small spacecraft flights that incrementally advance capability while building the heritage and experience of the technology and the team.

To encourage and develop understanding, it's in no small part a classic case of technologists needing to listen to those who might someday use their technology. We must listen to learn what the scientists' and mission-managers' pinch-points are. Only then can we endeavor to explain how this technology may be of service to the scientific goals.

If the autonomous systems community can successfully do these things, we as a society stand to enter an exciting new period of human and robotic discovery.

References

Boeing Integrated Defense Systems. Orbital express demonstration system overview. Technical report, The Boeing Company, 2006. Available as http://www.boeing.com/bds/phantom_works/orbital/mission_book.html

R. Castaño, T. Estlin, D. Gaines, A. Castaño, C. Chouinard, B. Bornstein, R. Anderson, S. Chien, A. Fukunaga, and M. Judd. Opportunistic rover science: finding and reacting to rocks, clouds and dust devils. In *IEEE Aerospace Conference*, 2006.

S. Chien, R. Sherwood, D. Tran, B. Cichy, G. Rabideau, R. Castano, A. Davies, D. Mandl, S. Frye, B. Trout, et al. Using autonomy flight software to improve science return on Earth Observing One. *Journal of Aerospace Computing, Information, and Communication*, 2(4):196–216, 2005.

Steve Chien, Richard Doyle, Ashley Davies, Ari Jónsson, and Ralph Lorenz. The future of ai in space. *IEEE Intelligent Systems*, (1541-1672/06):64–69, July/August 2006. Available as www.computer.org/intelligent

Richard J. Doyle. Autonomy needs and trends in deep space exploration. In *AVT Course on "Intelligent Systems for Aeronautics"*, number RTO-EN-022. RTO, May 2002.

T. Estlin, R. Castano, D. Gaines, B. Bornstein, M. Judd, R.C. Anderson, and I. Nesnas. Supporting increased autonomy for a Mars Rover. In *9th International Symposium on Artificial Intelligence, Robotics and Space*, Los Angeles, CA, February 2008. Jet Propulsion Laboratory, National Aeronautics and Space Administration.

T. Fong, M. Allan, X. Bouyssounouse, M.G. Bualat, M.C. Deans, L. Edwards, L. Flückiger, L. Keely, S.Y. Lee, D. Lees, et al. Robotic site survey at Haughton Crater. In *Proc. of the 9th Intl. Symp. on Artificial Intelligence, Robotics and Automation in Space*, Los Angeles, CA, 2008.

Jeremy Frank. Automation for operations. In *AIAA Space Conference and Exposition*. Ames Research Center, National Aeronautics and Space Administration, 2008a.

Jeremy Frank. Cost benefits of automation for surface operations: Preliminary results. In *Proceedings of the AIAA Space Conference and Exposition*. Ames Research Center, National Aeronautics and Space Administration, 2008b.

S.C. Hayden, A.J. Sweet, S.E. Christa, D. Tran, and S. Shulman. Advanced Diagnostic System on Earth Observing One. In *Proceedings of AIAA Space 2004 Conference and Exhibit, San Diego, California, Sep. 28-30, 2004*, 2004.

R. Knight. Automated planning and scheduling for orbital express. In *9th International Symposium on Artificial Intelligence, Robotics, and Automation in Space*, Los Angeles, California, February 2008. Jet Propulsion Laboratory, National Aeronautics and Space Administration.

Marvin Minsky. Telepresence. *OMNI magazine*, June 1980. Available as http://web.media.mit.edu/ minsky/papers/Telepresence.html.

Nicola Muscettola, P. Pandurang Nayak, Barney Pell, and Brian C. Williams. Remote agent: to boldly go where no ai system has gone before. *Artif. Intell.*, 103(1-2):5–47, 1998. ISSN 0004-3702. doi: http://dx.doi.org/10.1016/S0004-3702(98)00068-X.

A. Ogilvie, J. Allport, M. Hannah, and J. Lymer. Autonomous Satellite Servicing Using the Orbital Express Demonstration Manipulator System. In *Proc. of the 9th International Symposium on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS'08)*, pages 25–29, 2008.

B. Pell, D.E. Bernard, S.A. Chien, E. Gat, N. Muscettola, P.P. Nayak, M.D. Wagner, and B.C. Williams. A remote agent prototype for spacecraft autonomy. In *SPIE Proceedings Vol. 2810-Space Sciencecraft Control and Tracking in the New Millenium*, Denver, Colorado, August 1996.

Jonathan V. Post and Donald D. Rose. A.I. In Space: Past, Present & Possible Futures. URL http://www.magicdragon.com/ComputerFutures/SpacePublications/AI_in_Space.html.

B.C. Williams and P.P. Nayak. A model-based approach to reactive self-configuring systems. In *Proceedings of the National Conference on Artificial Intelligence*, pages 971–978, 1996.