# Autonomous Rover Technologies

**The Robotics Institute, Carnegie Mellon University**
**February 14, 1994**

**Submitted to: National Aeronautics and Space Administration**
**For the period: October 1, 1994 to September 30, 1995**

**Principal Investigators:**
**Eric Krotkov, Reid Simmons, William Whittaker**

## 1. Objectives

This project is developing innovative technologies in the areas of rover configuration, 3D perception, and planning and task-level control. Recent developments include design and fabrication of a novel frame walker, robust stereo perception and mapping software, and investigations into modeling and planning techniques that take uncertainty and risk into account.

The Autonomous Rover Technologies project will continue to perform fundamental, basic research that is critical for robotic exploration of planetary surfaces. In the area of robot configuration, technologies will be developed to aid in designing and characterizing robot performance, to formulate tradeoffs in different designs, and to perform high-fidelity simulations of configurations in order to evaluate and verify performance.

Perception technologies will be developed in the areas of stereo mapping using weak calibration methods, position estimation based on map registration and visual landmarks, and large-scale mapping based on fractal interpolation and topographic analysis. These technologies support robots that operate with uncertain sensing and that travel over large distances.

Task-level control technologies will focus on planning in the face of risk and uncertainty, two hallmarks of exploration missions. The objective is to make robots aware of their own capabilities and limitations, and to utilize that knowledge to keep themselves safe while achieving complex tasks. We are developing novel probabilistic and decision-theoretic planning techniques for this effort, with the expected results having both theoretical and practical impact.

## 2. Justification

Robots that operate for extended periods of time in natural, unstructured, and potentially hazardous environments need significant abilities to sense, plan and act in ways that maintain their safety and integrity, while achieving the high-level tasks commanded by remote human operators. The basic research undertaken in this project focuses on developing autonomous robot competence for self-reliant operation in rugged, uncertain environments. The project is developing both fundamental insights and practical tools for maximizing robot autonomy and reliability, of the type needed for planetary exploration.

The technologies being developed have been chosen with regard to their relevance to remote robotic planetary exploration. Such missions will inevitably require the utmost in robot reliability and flexibility to respond to unexpected and unpredictable events. In particular, the results of this project will feed into and directly respond to the Lunar Rover Demonstration focused research thrust.

This research program will have implications far beyond the achievement of any specific task. The program will develop scientists, insights, and technologies that are crucial to the realization of NASA's long-term objectives. The work leverages our current efforts in building and controlling autonomous robots, and dovetails with the activities of NASA centers engaged in robotic research.

## 3. Approach

The basic research covered by this proposal falls into three broad categories: rover configuration, perception, and planning and task-level control.

### 3.1 Rover Configuration

Mobile robot design combines diverse technologies exhibiting intricate interdependencies. The effects of subsystem parameters and configuration changes on overall performance are subtle enough to defy intuition and simplistic modeling. Despite the difficulty, evaluations must be made to create designs, and this currently occurs with costly and imperfect mixes of expertise and intuition. Unlike manipulator counterparts, mobile robot configuration is a fledgling and fast moving discipline that does not yet benefit from deep experience, static knowledge or generalized solutions.

Most mobile robot designs have emphasized certain functionalities while ignoring or being oblivious to others. Hence the proliferation of imbalanced or dysfunctional prototypes that fall short of explicit or implicit objectives. The current state of configuration is at best a satisficing process. The need for competent robots gives rise to the need for tools and processes that will yield more complete and optimal robots. Appropriate tools will incorporate expertise of robot subsystem functionality and will automate the systemic tradeoff and optimizing processes.

We propose to develop a framework that will assist a designer in making tradeoff decisions. This will enable a robot designer to more easily pursue subdesigns that have succeeded previously or to abandon those and attempt more global searches for successful design combinations. Precedents of prior robot design processes will be incorporated as a foundation for this framework. The framework will sew together models of varying specificity and realism to provide some solution prospects despite missing or estimated information in early design stages.

We will develop subsystem models to populate our framework. Initially, the work will provide design models for wheeled vehicles in realistic terrain with power, thermal, communication, sensing and planning means for exploratory navigation. These tools will allow designers to explore and validate ideas with consideration of the important system interdependencies. The tools will also allow functional simulation of missions and exploration of component failures.

Previous experience has proven the benefit of simulators for evaluating new ideas, identifying pitfalls, and as virtual vehicles for validating and tuning navigation software. The proposed software will go beyond simulation and analysis to help design more capable, reliable mobile robots and to speed up the development process.

Our approach will be to build from existing simulation and analysis tools which are a mix of research and commercial building blocks. We will create the integrating framework, incorporate component models, then refine the subsystem models to function in the absence of complete knowledge. We will apply this framework to the design of the Lunar Demonstration Robot and evaluate the utility of the framework as a design and simulation tool.

### 3.2 Perception

Perception interprets sensor signals, relating them to the state and structure of the environment. Interpretation of sensor signals plays essential roles for planetary rovers traveling in unknown territory, enabling safe traversal around local obstacles, on-board estimation of position, and incremental construction of models of the terrain covered by the robot.

To advance perceptual capabilities for rovers, the proposed research will concentrate on three central topics:

1) Mapping local surface geometry using stereo vision,
2) Estimating rover position using multiple sensors, and
3) Mapping large-scale surface geometry.

The research will generate fundamental insights and innovative perceptual techniques that will feed into focused efforts to traverse hundreds of kilometers on natural terrain.

#### Mapping Local Surface Geometry

Mapping the local surface geometry is critical for safe and power-efficient locomotion. We are investigating two approaches to acquiring and exploiting three-dimensional stereo information, one model-intensive, the other embedded. The model-intensive approach builds elevation maps from stereo data; these maps can be used for detailed geometric path planning. The embedded approach embeds stereo into the planning and control system without explicitly computing maps; this forms the basis of a fast system for autonomous driving.

In FY94, we successfully demonstrated stereo mapping using weak camera calibration. In FY95, we will develop and demonstrate a fast stereo obstacle avoidance system.

Specifically, in FY95 we will develop an embedded stereo driving system, and demonstrate the system operating a research vehicle. We will divide the effort into two principal parts. In the first, we will develop methods for evaluating the traversability of terrain given weakly calibrated stereo images. The traversability will be determined by shape indicators such as relative slope and relative elevation (since these quantities do not require explicit scene reconstruction or mapping, it is this step that provides the speed advantage over traditional methods). In the second part, we will develop algorithms for evaluating paths based on traversability images, and for computing new driving commands for the vehicle based on the path evaluations.

#### Estimating Rover Position

Estimating position and attitude is a fundamental capability required for many mobile robot tasks, and is critical for long-range traverses and accurate reconnaissance on planetary surfaces. We are investigating a family of approaches, including map/map registration, visual landmark alignment, and multi-sensor estimation.

In FY94, we successfully attacked the map/map registration problem using geometric hashing (efficient) instead of search (computationally intractable). In FY95 we will concentrate on a maximum likelihood estimator for triangulation on visual landmarks.

Specifically, we will address the visual landmark triangulation problem in the presence of outlier measurements by applying techniques from tracking and from statistical hypothesis testing to develop rigorous and robust solutions. The research will proceed in the following stages.

1) Formulate, implement, and simulate a maximum likelihood estimator with different and explicit stochastic models for small errors and large outliers.

2) Evaluate the performance of this estimator with synthetic images, for example images of the Rocky Mountains generated by the “maptool” package from the University of Utah.
3) Develop feature detectors suitable for rugged, high-relief environments.
4) Demonstrate maximum likelihood position estimation using features detected in real images of natural, high-relief scenes.
5) Compare performance of conventional Kalman filter and the robust maximum likelihood estimator.

### Mapping Large-Scale Surface Geometry

In previous work, we developed methods for building topographic maps by identifying features such as peaks, pits, ridges, and ravines from elevation maps of reconstructed surfaces. We will extend this line of analysis by developing more rigorous methods of computing iso-elevation contours, developing a more stable method for extracting high-curvature points, and devising better strategies for linking ridge (ravine) points together. These topographic features will be valuable inputs to route planners, whether automated or human.

We will incorporate the topographic analysis techniques into a single package. This package will read in elevation data and digital line graph data from diverse sources, including the United States Geological Survey, Defense Mapping Agency, and selected perception systems. The package will interactively render the data and the results of the topographic analysis techniques using overlay planes. We will demonstrate the utility of the system by providing geographic information on the Lunar Rover Demonstration sites, such as the Pittsburgh slag dump or the White Sands gypsum fields.

## 3.3 Planning and Task-Level Control

Reliability is of utmost importance for both autonomous and telerobotic planetary missions, due to the spatial and temporal isolation from ground-based control. Robot systems must be designed with reliability in mind, and the robot systems themselves should have an understanding of their capabilities and limitations in order to maintain their self-integrity (what we call ‘‘cognitive reliability’’). We propose to pursue fundamental issues in robot cognitive reliability, with the goal of enabling robot systems to survive for extended periods of time in natural, unstructured, remote environments.

We are investigating the issues of reliability in several areas of the planning and task-control problem. Their relationships can be seen in the figure below: reliable models of the environment are needed for planning, and planning processes are controlled by task-level control strategies. The research efforts described below will combine to produce overall reliable planning and control systems.

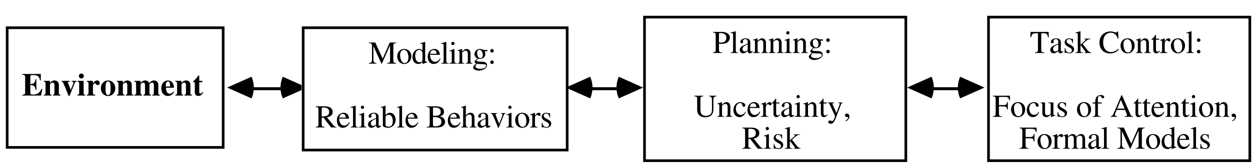


Uncertainty is a hallmark of acting in unstructured environments: one is uncertain what the environment will be like, about the effects of actions, or even how the robot itself will behave. It is our contention that the reliability of mobile robot systems can be significantly increased by explicitly taking this uncertainty into account when generating plans. We are investigating several approaches that utilize uncertainty in the form of probabilistic and decision-theoretic reasoning to produce more reliable robot systems.

### Modeling Reliable Behaviors

While probabilistic information is well-suited for reasoning about uncertainty and risk, its utility is highly dependent on having good probability estimates. While one can estimate the probabilities through experimentation in the environment (or analogs), the results are never perfect. In particular, there may be unmodeled interactions between behaviors that produce unintended effects.

We are developing a novel formalism that uses lower probabilities (akin to Dempster-Shaffer belief functions) to represent the limitations in action models. The formalism can be used to reason about how robust a particular behavior is, that is, what is the range of probabilities that the behavior will have its intended effects.

We intend to use this formalism to develop techniques that assist robot system designers in developing reliable behaviors. The designer would specify a state space as detailed as current knowledge allows and design conditional, iterative behaviors using those action models. The effects of the behaviors would be projected using the lower probability model: if the range of the result is small, this indicates that the behavior is highly predictable (even given the uncertainties within the model) and thus the behavior is likely to be very reliable in practice. If the behavior is found not to be highly predictable, a sensitivity analysis of the behavior can reveal places where more precise knowledge about the environment would improve the prediction. Thus, these techniques can be used to help a designer identify those conditional behaviors that are highly predictable, and thus reliable, versus those with a weak spot where an unmodeled factor could cause troubles.

The main goals of this line of research are to develop the novel formalism for reasoning about lower probabilities, including methods for automatically projecting and analyzing conditional behaviors. These general methods will be applied to the planetary exploration domain as tools to aid designers in generating highly reliable plans of action.

### Planning Under Uncertainty and Focus of Attention

In planning, one wants to produce the best possible plan without expending undue computational resources. This *meta-planning* problem of deciding when to plan and when to begin acting is especially important when dealing with uncertainty.

We are developing techniques for planning under uncertainty that perform a sensitivity analysis on plans to determine where to focus planning effort. The techniques can determine if one plan is clearly dominated by another (and thus should no longer be considered); if the uncertainties in a plan are high enough, the sensitivity analysis can pinpoint where to expend more planning effort — either by planning at a higher level of detail (such as by using a more detailed models, or more complete planners), or by planning to perform sensing operations to reduce uncertainties. The techniques use probability intervals, such as those described in the modeling section above.

The basic idea is that in environments where the uncertainty is low, or the difference in the utility of alternate plans (the opportunity cost) is low, the system will use minimal planning effort. As the uncertainty gets higher, or the opportunity costs of different plans increase, the planner will expend more effort, focusing on the critical differences between competing plans. The emergent effect will thus be faster navigation in open, well-known, spaces and careful planning, sensing and execution in situations where the cost of making mistakes is high. We plan to test these ideas in an analog planetary rover mission, consisting of extended navigation and exploration tasks.

### Planning Under Risk

Traditional approaches to planning take a risk-neutral approach. This is clearly undesirable for planetary rovers: in fact, we would like to be able to modulate the risk attitude of rovers: making them risk aversive, at least initially, possibly becoming more risk-seeking as missions progress and confidence in their abilities increases.

We are investigating extensions to probabilisitic planning methods to explicitly include a factor for risk. We have already shown, for instance, how to transform certain risk-sensitive planning problems into pure probabilistic planning problems. We are beginning to study how to learn the risk attitude of other agents. This may be useful in interactive mission scenarios, where remote humans are commanding the rover. By knowing their risk attitude, the rover can optimize behavior, such as by tailoring the sensor inputs the operator is likely to need, or by monitoring the situation more closely if the operator is found to be risk-seeking. We are also investigating the tradeoffs between risk and other factors, such as time, power utilization, etc.

Note that many of these planning techniques are computationally quite expensive. We therefore envision that they will be used mainly as tools to help designers develop plans and understand the consequences of decisions, rather than being used directly on the rover in a run-time environment. Given suitable computational resources, however, it is desirable to put some of this reasoning on board, so that the rovers can reason about and react to novel situations in an intelligent fashion.

### Formalization of Task-Level Control Systems

In FY94 we have started to produce a formal model of TCA (using the Z notation and temporal logic). This formal model will be used as the basis for verification and analysis tools described in the “Task Control and Communications” proposal to TRIWG. In FY95, we will complete the formalization, adding models for the monitoring and exception handling control constructs of TCA. This task is being performed under this proposal, because it is very much a basic research endeavor, necessitating the utilization and development of novel techniques for modeling and reasoning about concurrent, distributed systems.

## 4. Participants

The proposed work will be undertaken at the Robotics Institute at Carnegie Mellon University. The work will span several labs and centers within the Robotics Institute, including the Field Robotics Center, the Vision and Autonomous Systems Center, and the Learning Robot Laboratory. Participants in the research include faculty members, graduate students, and undergraduate students.

## 5. Deliverables

The following include the deliverables for this project in FY95.

Configuration:

- Demonstrate a framework and its automated tradeoff processes as applied to Lunar Rover configuration and two additional non-lunar robots. Check performance on the non-lunar robots against formerly completed designs.
- Demonstrate the high fidelity simulator tools by showing three different configurations performing a lunar mission and the evaluated differences in performance.
- Demonstrate the validity of the generated component and subsystem models by verifying them against those in existing robots. Compare the model behavior in lunar conditions against specific, detailed analyses and lunar experience.

Perception:

- Demonstrate an embedded stereo obstacle avoidance system that computes driving commands 10 times faster than a traditional stereo approach that generates driving commands based on analysis of elevation maps computed from stereo disparities.
- Demonstrate that maximum likelihood estimator using visual landmarks can handle twice as many outliers as Kalman filter using same visual landmarks.
- Implement large-scale mapping package that computes four topographic features not available in current software systems.

Planning and Task-Level Control

- Produce a complete formalization of the TCA control constructs, including task decomposition, resource scheduling, monitoring and exception handling. Transfer this technology to be used in the "Task Control and Communications" project.
- Demonstrate techniques for reasoning about the robustness of behavioral sequences. Use to quantify the reliability of the basic action sequences of the Ratler rover.
- Demonstrate the use of sensitivity analysis in focusing planning effort in probabilistic, uncertain situations. Quantify increase in the plan quality versus savings in planning effort by comparing experimentally to classical planners that consider all possible plans and reactive systems that act without planning.

## 6. Commercialization Potential

Given the long lead times in a basic research project, it is typically quite difficult to accurately gauge the commercial potential of the anticipated results. Often, the major impact of such work is in the papers and articles that describe the research. Nevertheless, it is clear that the results developed in this project can have significant commercial impact in applications that are characterized by robots operating remotely in unstructured, natural, and potentially hazardous environments. In particular, we expect that applications of the technologies developed here will impact mining, agriculture, construction, and toxic waste remediation.

Even if such robots are not completely autonomous, the *capability* to act autonomously means that the robots can be more self-reliant, and can detect and avoid situations that might otherwise prove injurious. By making them safer and more effective, we both increase the perceived value of the robots, because they can be used in more environments and can tackle more variety of tasks, and decrease their overall life-cycle costs, because they can survive for longer periods of time. The cumulative effect is to make mobile robotics a more attractive venue for utilization and investment.

## 7. Funding

The requested funding for the Autonomous Rover Technologies program for FY95 is $520,000. This covers partial support for the principal investigators, partial support of a research scientist, six graduate students, travel to refereed conferences to present the results of the research, and a minimal amount of equipment money, mainly related to configuration and perception efforts.

| | FY95 |
|---|---|
| Personnel & Overhead | $500,000 |
| Equipment | $20,000 |
| | |
| Total for FY95 | $520,000 |

## 8. Facilities

The Robotics Institute at Carnegie Mellon University is a unique resource. Established in 1979 as the nation's first comprehensive robotics research center to conduct research in robotic technologies relevant to scientific, industrial, and governmental needs, the Institute has grown into a powerful, interdisciplinary research organization, which draws upon Carnegie Mellon's strengths in computer science, engineering, science, and industrial administration.

The Institute has a staff of over 120 personnel, over 50 graduate students enrolled in affiliated programs, and 60 graduate students enrolled in the Robotics Ph.D. program. The Institute is organized into thirteen dedicated research laboratories and three program centers, which cover the spectrum of robotic research initiatives and encompass a wealth of robotic expertise. The Institute pursues a comprehensive research agenda that culminates in competent, working robots.

## 9. Principal Investigators

Eric Krotkov is a Research Scientist in the Robotics Institute at Carnegie Mellon University. Dr. Krotkov has been a faculty member at CMU since 1988, and prior to that worked at LAAS in Toulouse, France.
(phone: 412-268-3058; fax: 412-621-1970; e-mail: epk@cs.cmu.edu)

Reid Simmons is a Research Computer Scientist in the School of Computer Science and Robotics Institute at Carnegie Mellon University. Dr. Simmons has been a faculty member at CMU since 1988.
(phone: 412-268-2621; fax: 412-621-5739; e-mail: reids@cs.cmu.edu)

William "Red" Whittaker is a Principal Research Scientist in the Robotics Institute at Carnegie Mellon University. He is founder and director of the Field Robotics Center at CMU.
(phone: 412-268-6559; fax: 412-682-1793; e-mail: red@frc2.frc.ri.cmu.edu)

The mail address of the three principal investigators is:
Robotics Institute
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213