A Rule-based Strategy for Astronaut Following Operations

Jamie A. Lennon[†] **Ella M. Atkins**[‡]
University of Maryland, Space Systems Laboratory, Building 382
College Park, MD 20742
{jlennon, ella}@ssl.umd.edu

Abstract

Scenarios for a manned mission to Mars call for astronaut extravehicular teams to be accompanied by semi-autonomous rovers. These rovers must be able to safely follow the astronauts over a variety of terrain with minimal assistance. We propose a color-based visual tracking system as a high-speed sensory approach for astronaut following. In previous work we characterized the accuracy of the system, which was found to be highly dependent on environmental conditions. In this paper, we propose a set of rules to enable safe following behavior, concentrating on reactions when the visual target cannot be clearly identified.

Introduction

When humans go to the moon and Mars, they will undoubtedly take robots with them. Robotic assistants can greatly enhance crew productivity, either by autonomously completing tasks on their own or assisting the astronauts in some way (Dorais et al. 1998). Astronauts on surface extravehicular activities (SEVAs) will certainly find rover assistants to be useful. A rover can carry cumbersome tools and heavy samples. It can also provide an additional measure of safety if it carries extra life support and communications equipment.

Such a rover will ideally have a multitude of operational modes. The astronauts may wish to ride on it and drive to a remote site. They may desire to direct the rover to autonomously approach a marker or beacon and perform some arduous task such as core sample extraction. The rover could be teleoperated by a crewmember at the base camp. At least some of the time, we expect that the SEVA crew will be walking over the Martian surface and will want the rover to follow them. For instance, if astronauts ride the rover to a remote site to take samples, they may wish to disembark to look for interesting rocks. It will be very useful if the rover can follow them about, keeping a variety of tools handy, taking detailed *in situ* data on the samples before their removal, and then accepting the samples from the astronaut.

In this paper, we describe a modal operations model based on a color-based tracking system to enable such behavior. In previous work (Lennon and Atkins 2001), we tested the Cognachrome^{TM}-based tracker on a simulator that emulates the changes in relative position and velocity between the astronaut and rover and assessed its "heeling" and “following” accuracy. As described in (Atkins, Lennon, and Peasco 2002), we also have demonstrated target following with the University of Maryland Space System Laboratory's free-flying Secondary Camera and Maneuvering Platform Space Simulation Vehicle (SCAMP SSV).

A number of factors affect ability to accurately track color-blob targets. First, lighting must be relatively constant and the tracked object should be minimally reflective. Particulate (dust) can block objects, and terrain (e.g., large rocks) can block the target view. Therefore, alternate navigation schemes must be available. This paper presents rules by which high-speed color tracking, low-speed inertial navigation, and teleoperation/manual control can be coupled to form an efficient, robust rover companion. Work is in progress to implement this multi-mode system within a practical robotic system.

Background

The Mars Reference Mission defines at least three types of rovers needed for planetary exploration (Hoffman and Kaplan 1997). The first class is for use in the near vicinity of the base camp and may be nothing more than wagons or carts. The second class is the unpressurized rover similar to the Apollo lunar rover that would be used for a six to eight hour sortie away from base. The third class is the mobile pressurized habitat, allowing crew to undertake ten-day missions far from the base. While this third sort is to have robotic arms for sample collection, it also has an airlock, so that astronauts may disembark. It may be desirable for both rovers of the second and third class to follow a mobile astronaut over the planetary surface. It is *not* desirable to occupy a SEVA astronaut or an astronaut at base with any more rover operations than necessary (Washington et al. 1999). We would like for the rover to be able to follow the astronaut on its own.

NASA experiments using the Marsokhod rover are testing the "follow the astronaut" paradigm (Kosmo, Trevino, and Ross 1999). These tests show that the rover is useful only if it can keep up with the astronaut. Otherwise, the astronaut must spend a great deal of precious SEVA time waiting for the rover to catch up. This means the rover needs to travel at speeds of about 1.4

[†] Graduate Student Fellow, Aerospace Engineering Dept., University of Maryland
[‡] Assistant Professor, Aerospace Engineering Dept., University of Maryland

m/s. (This is an average walking speed for an unencumbered human on earth, and it is unlikely that a suited astronaut will exceed it.) Previous rover research testbeds combined teleoperation with autonomous navigation and obstacle avoidance using stereo vision. Such systems could achieve speeds of around 0.5 m/s (Simmons et al. 1995), too slow for following an astronaut at a "natural" gait.

An astronaut in SEVA is already performing navigation and some obstacle avoidance as he or she traverses the Martian terrain. We wish to take advantage of advanced human perception capabilities to augment rover navigation. In "following" mode, the rover does not need to perform extensive analyses of the surrounding terrain if it can rely on its human partner to pick a safe path. It needs only to record that path and then follow it while maintaining a safe distance between itself and the astronaut. How will it observe and record that path?

Efficient resource utilization – including computing resources – is one of the essential components of a successful autonomous system (Washington et al. 1999). Color blob tracking is a simple and fast vision process. It has been proven to be fast, accurate, and robust in several robotics competitions (Bailey et al. 1998) and also does not require the astronaut to carry any emitters or equipment other than a small colored target. While a gesture-based system would allow for more complex communications between the rover and astronaut, it would also require either more complicated vision processing or transmission of astronaut finger joint angles to the rover. Note that although we assume color blob tracking in the below discussion, an analogous following strategy would apply for alternative relative navigation signals subject to loss.

Modal Operations

Color blob tracking is useful in a variety of situations. But there are many situations where the rover could find lighting conditions impossible to handle or simply lose sight of the astronaut. How shall it proceed in these all too-frequent cases?

A Biological Paradigm

A possible answer lies in examining another multiagent expeditionary force - the family on vacation. Generally, a few members (the adults) fulfill the role of the astronaut, selecting areas of interest and leading the rest of the family to them. The children tag along behind, through complicated fields of mobile obstacles that bear a strong resemblance to their tracking targets. To assist their children in tracking them, many parents provide extra visual clues - a balloon of a certain color high overhead, or a distinctive hat worn on a tall adult. But as we all know, even these augmented color blob tracking methods can fail, and the child may find himself separated from the adult.

What the child does in this situation depends on how he has lost track of the adult and where he is at the moment. If the child cannot see the adult because the adult just walked behind the corner of a building, the child if sure that if he also continues around the corner of the building, he will see the adult again. This loss of signal is far less troubling than when the child, preoccupied by something else, looks around to discover that he has lost sight of the tracked object. He does not know where the adult has gotten off to, in what direction, or how far away he might be. If the child feels truly lost, he will engage in a behavior to aid in becoming found again. The behavior selected may depend on the environment.

A child in an unknown, potentially hostile environment has probably been instructed to do one of two things. First, he can seek out a known safe spot or landmark - a police officer or a meeting place established in case of separation. Second, he can stay where he is so his parents can backtrack and find him. Broadcasting a distress signal is an option in both cases.

A child in a known or structured environment might feel comfortable searching for his parents. If separated in, for instance, a grocery store, a simple pattern search checking each of the aisles is likely to reveal the location of the parents. The child can use his knowledge of his parents' shopping patterns (starting at one end of the store and going through the aisles, moving slowly in the aisles as they shop) and his other beliefs (that they could also be in line at the checkout, that they won't leave without him) to aid in the search. A simple visual search can again be augmented by other means, particularly widely broadcasting a signal of some kind. The child might call out for his parents as he looks down the aisles, or might even ask the store owner to use the public address system to call them.

Applying the Paradigm to the Rover

A rover can be programmed with all of these behaviors. For this discussion, we presume a rover with camera mounted on a PTU (pan-tilt unit) is used to perform color-blob tracking. We presume the operator has set the rover to prefer “following” mode to maximize travel speed and minimize operator overhead. Figures 1 - 3 show the schematics for recovering from situations in which color tracking capabilities are potentially compromised. We present recovery strategies both for a rover losing signal while tracking and the case when it disengages from some other task and cannot find the astronaut/target. In all figures, encountering a “Follow” block implies high-speed color blob tracking may again be engaged, while a “Stop” block implies human intervention potentially in the form of inertial navigation commands or direct teleoperation is required.

When the tracked object is lost, the rover should first consider how it was lost. First, the hardware is checked for nominal performance. If either camera or Cognachrome™ is not working, tracking cannot occur. Figures 1 and 2 consider the cases where the target was being tracked when it was lost. Figure 3 considers the case where the rover “looks up” from another task and cannot see the astronaut.

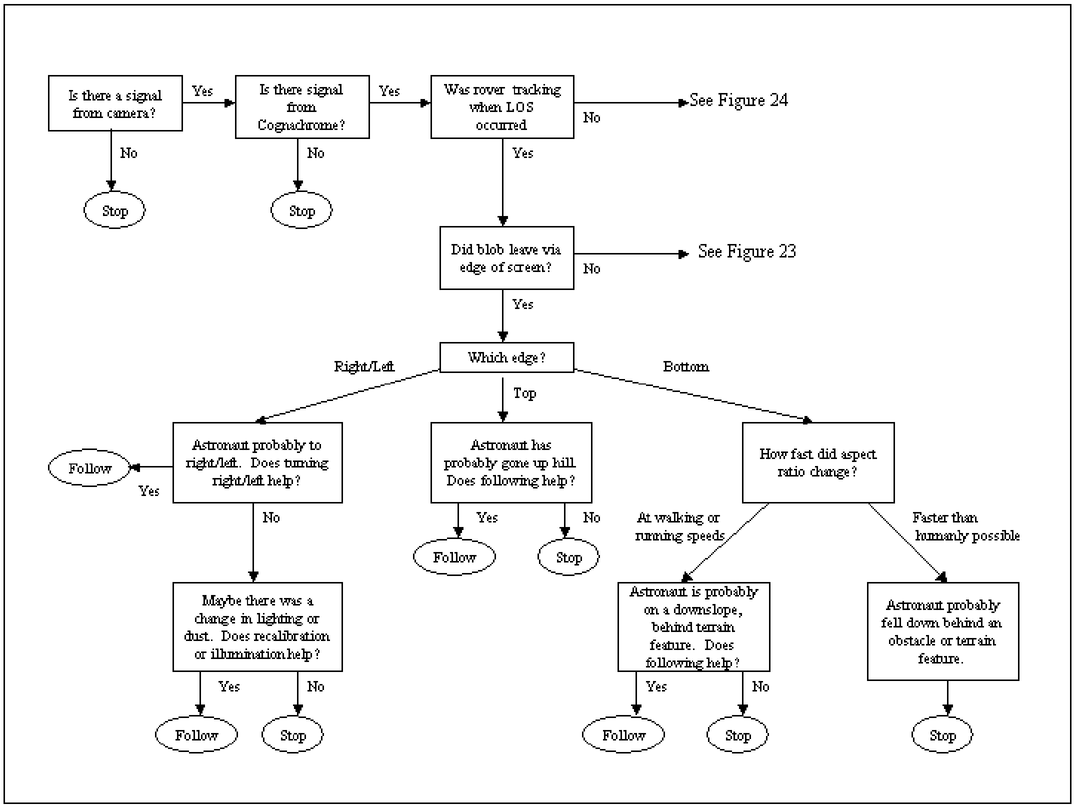

Figure 1: Reacquiring signal when target exits edge of screen

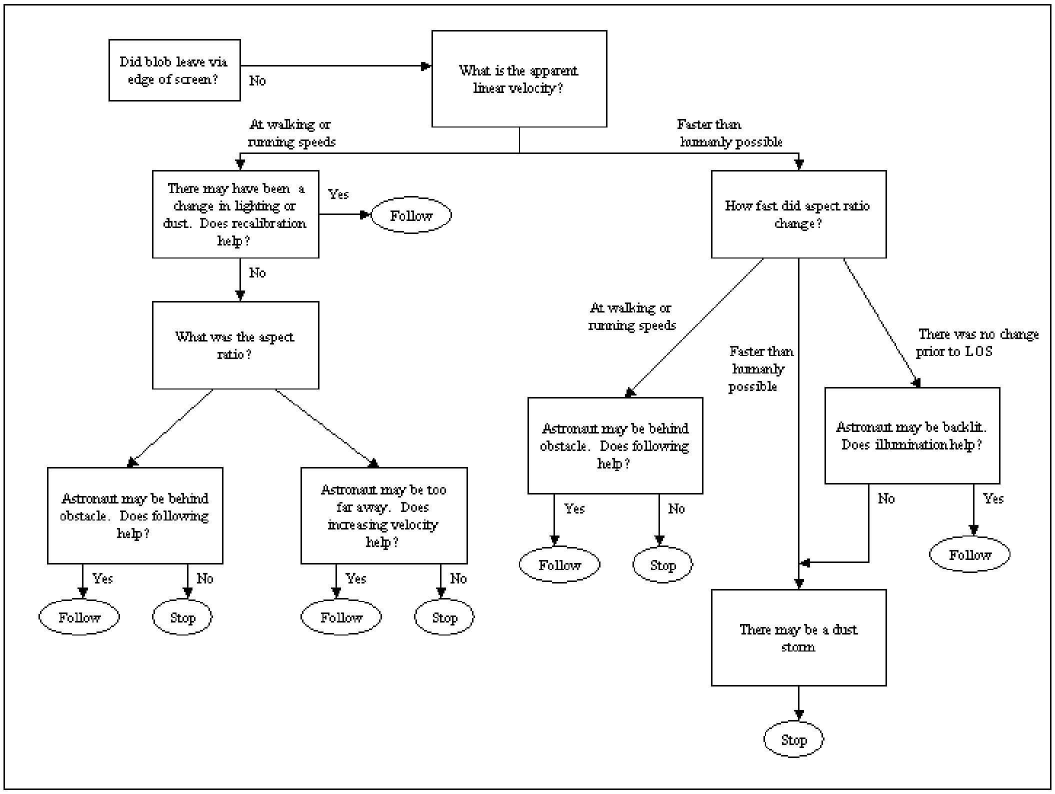

Figure 2: Reacquiring signal when target disappears from near-center screen

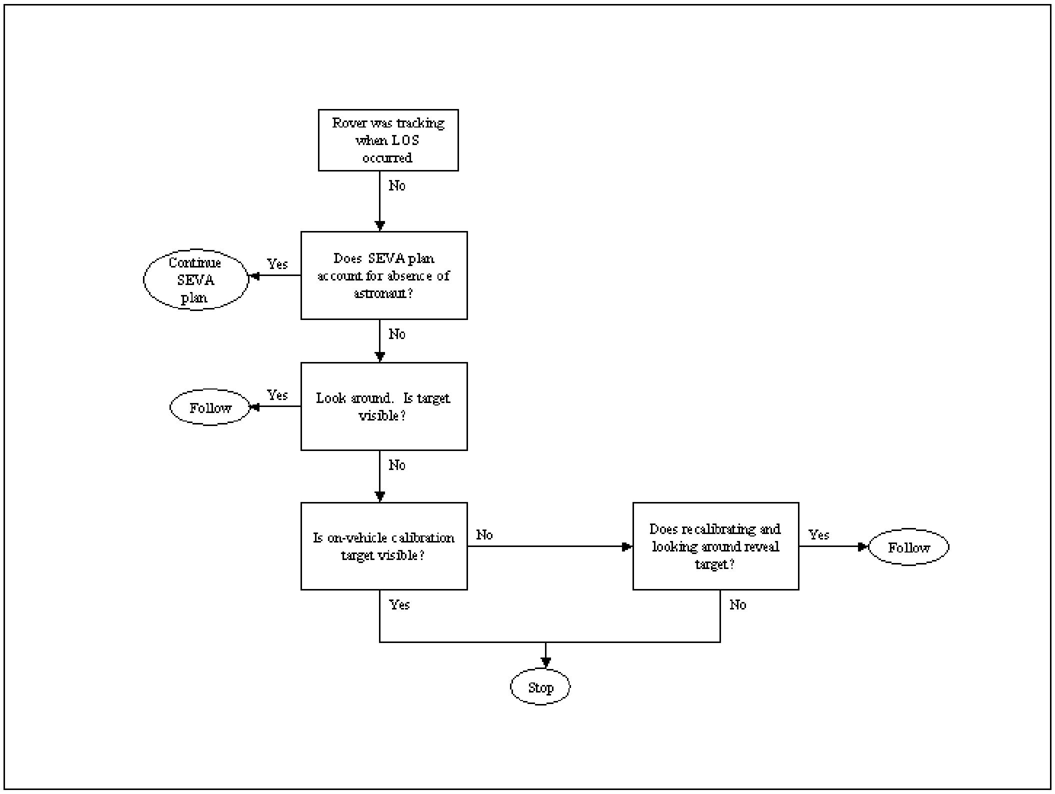

Figure 3: Reacquiring signal when target disappears while rover is not tracking it

Figure 1 considers those cases where the astronaut has left the field of view under his own power. These are the cases when the target is moving left to right, or up to down, too quickly for the PTU to track (or, if the PTU has reached a hard stop, for the rover to turn) and goes off-screen. They can be identified by examining the centroid data prior to loss of signal (LOS). The last known centroid position can tell us whether the target was at the edge of the screen. If the centroid's velocity is also towards that edge, we may say that the centroid moved off-screen. If the velocity was in a different direction or was zero, then the target did not move of its own accord off-screen, and the LOS is probably the result of one of the causes found in Figure 2.

For clarity in Figure 1, two separate procedures were combined in the questions, "Does turning right/left help?" and "Does following help?" In the case of turning, the PTU should obviously turn until it reaches a hard stop before turning the rover. It simply takes less power to move the PTU than the entire vehicle. For the top/bottom following, the PTU should first tilt in that direction until a hard stop is reached, to see if this brings the astronaut back into view. If that does not, however, the rover should carefully embark upon the astronaut's path up to the point of disappearance. We know that the astronaut cannot fly, so we assume that any disappearance off of the top of the screen is the result of the astronaut climbing a terrain feature. As the rover physically follows, it too will ascend the terrain feature, and the astronaut should become visible again. The case of the bottom of the screen must be more cautious, since there may be holes or rocks into or behind which the astronaut may fall. The case of falling is distinguished from the case of the astronaut descending a terrain feature by the rate at which the aspect ratio of the target changes as it goes off-screen. We assume that an astronaut falling uncontrollably will leave the screen much more quickly than one walking over the crest of a hill. The rate of change in the aspect ratio of the target as it leaves the screen will then also be greater for a fall than for walking.

Figure 2 covers those cases where the target does not leave the screen under its own power. Changes in lighting and dust accumulation can affect Cognachrome™ calibration to the point where the target is no longer visible. A sudden, violent Martian dust storm may obscure all vision. Or it may be that the astronaut has simply stepped behind a rock, or has gotten too far away.

Since the area of the target is used to calculate the range from the rover to the target, the change in area can be used to compute the linear velocity of the target – assuming the entire target is visible. If the astronaut were to step behind a rock, the sudden decrease in target area would indicate that he was dashing across the landscape at high speeds! We can use our knowledge of the astronaut's range of possible velocities to separate these cases. Of course, the astronaut may drift behind the rock very slowly. In these cases, we examine the aspect ratio just before LOS. If the aspect ratio was very close to 1, then the entire target was on-screen until LOS, and it is likely that the astronaut has simply gotten out of camera range. (This can, of course, be checked with the last known range of the astronaut). If the aspect ratio were not 1, then the target was occluded before LOS, and the astronaut has probably stepped behind some obstacle.

In the cases where the area decreased too suddenly for it to be the result of astronaut motion, we again turn to the rate of change in the aspect ratio to distinguish between cases. A change in aspect ratio conforming to astronaut movement reinforces the idea that the astronaut has stepped behind something. A very, very rapid change in aspect ratio could indicate that a dust storm has blown across the area. What if the aspect ratio did not change at all, but was constant up until LOS? The target may have suddenly become backlit (which is remedied with illumination). A dust storm that maintained target symmetry until LOS is another possible culprit.

Figure 3 handles the "child lost in the store" scenarios. It may be that the SEVA plans called for the rover to engage in some activities independent of the astronaut. If this is the case, then everything is going according to plan, and there is no cause for alarm. But what if the rover expects to see the astronaut in the area when it has finished its task? The first step, of course, is to look around for the astronaut. If the target is not immediately apparently, there is also the possibility that lighting or dust accumulation has changed the apparent color of the target. To check, one can paint a small patch of the target's color on the rover itself, somewhere where the PTU can precisely aim the camera at it and expect back a certain area, centroid, and aspect ratio. If this on-vehicle target is invisible, the rover can use it to recalibrate itself. (If the on-vehicle target is entirely obscured by dust, recalibration will obviously fail and the astronaut cannot be tracked, even if in the area). Now certain that its color information is updated to reflect the most recent environmental conditions, the rover can look around for the astronaut again. If the astronaut is still not visible, the rover is lost and waits for directions or assistance.

Conclusions and Future Work

In this paper, we have presented a rule-based behavior strategy for robust application of color blob tracking to the astronaut following problem, emphasizing recovery when the tracked target is lost. For a Mars rover already equipped with several different cameras, visual tracking is an attractive option. The position information gleaned from the camera(s) can be used in a control system to enable the rover to follow the astronaut over unknown terrain. In this way, we take advantage of the astronaut's superior navigation and planning abilities and free the rover from a time-consuming analysis of the local terrain and its hazards. Since the path the astronaut has taken is known to be safe, the rover can follow at velocities on the order of human walking speed, allowing it to keep pace with its human teammate. Should the rover lose sight of the human, either because of lighting conditions or because of other factors, mode-switching operations can enable it to determine its next best course of action.

We have tested basic following capabilities under "perfect" conditions in which the target can always be tracked. We have plans to test the spectrum of target loss situations in simulation and hardware environments, including adaptation of the recovery rules for application during free-flight and subsequent testing in a 3-D neutral buoyancy environment. Long-term, we intend to incorporate this knowledge in a planning/scheduling agent controlling high-level vehicle behaviors of which following is an important subset.

Acknowledgements

This material is based upon work supported under a National Science Foundation Graduate Research Fellowship.

References

Atkins, E., Lennon, J., Peasco, R. "Vision-based Following for Cooperative Astronaut-Robot Operations," *Proceedings of the IEEE Aerospace Conference*, Bozeman, Montana, March 2002.

Bailey, B., J. Reese, et al. "Robots with a Vision: Using the Cognachrome Vision System," *Circuit Cellar Ink*, Issue 92, March 1998.

Dorais, G., R.P. Bonasso, D. Kortenkamp, B. Pell, and D. Schreckenghost. "Adjustable Autonomy for Human-Centered Autonomous Systems on Mars," *Mars Society Conference*, August 1998.

Hoffman, S.J., and D.I. Kaplan, eds. *Human Exploration of Mars: The Reference Mission of the NASA Mars Exploration Study Team*. NASA Special Publication 6107, Johnson Space Center: Houston, TX, July 1997.

Kosmo, J., R. Trevino, and A. Ross. *Results and Findings of the Astronaut-Rover (ASRO) Remote Field Test Site*. JSC Doc. No. 39261, Johnson Space Center: Houston, TX, March 1999.

Lennon, J. and Atkins, E. "Color-based Vision Tracking for an Astronaut EVA Assist Vehicle," in *Proceedings of the SAE International Conference on Environmental Systems (ICES)*, Orlando, FL, July 2001.

Sherwood, R., Miskin, A., et al. "Autonomously Generating Operations Sequences for a Mars Rover using AI-based planning," *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, October 2001.

Simmons, R., E. Krotkov, et al. "Mixed-Mode Control of Navigation for a Lunar Rover," *Proceedings of the SSI/Princeton Space Manufacturing Conference*. 1995.

Washington, R., Golden, K., et al. "Autonomous Rovers for Mars Exploration," *Proceedings of the 1999 IEEE Aerospace Conference*, 1999.