Autonomous Robots 11, 341–346, 2001

Biologically Inspired Autonomous Rover Control*

TERRY HUNTSBERGER
Jet Propulsion Laboratory, MS 82-105, 4800 Oak Grove Drive, Pasadena, CA 91109, USA
Terry.Huntsberger@jpl.nasa.gov

Abstract. Robotic missions beyond 2013 will likely be precursors to a manned habitat deployment on Mars. Such missions require robust control systems for long duration activities. Current single rover missions will evolve into deployment of multiple, heterogeneous cooperating robotic colonies. This paper describes the map-making memory and action selection mechanism of BISMARC (Biologically Inspired System for Map-based Autonomous Rover Control) that is currently under development at the Jet Propulsion Laboratory in Pasadena, CA (Huntsberger and Rose, *Neutral Networks*, 11(7/8):1497–1510). BISMARC is an integrated control system for long duration missions involving robots performing cooperative tasks.

Keywords: biologically inspired control, action selection mechanisms, mobile robots

1. Introduction

Robotic outposts and precursor missions for deployment of manned habitat infrastructure on Mars are being studied by NASA for the second decade of this century (Schenker et al., 2000; Huntsberger et al., 2001). Such missions require more autonomy in their control architecture than the Pathfinder/Sojourner mission that landed on Mars in the summer of 1997. Representative biologically inspired navigation and control systems for such outposts include CEBOT (Fukuda and Kawauchi, 1993), Q-machines (Kube and Zhang, 1997), the Tropism System Cognitive Architecture (Agah and Bekey, 1997), behavior-based control (Mataric, 1997), ALLIANCE (Parker, 1998), and CAMPOUT (Pirjanian et al., 2000 and references therein). Good overviews can be found in Maes (1991), Pfeifer and Scheier (1999).

We have recently developed a multi-robot control architecture called BISMARC (Biologically Inspired System for Map-based Autonomous Rover Control) for long duration missions (Huntsberger and Rose, 1998). It is based on a free-flow hierarchy (FFH) similar to the DAMN architecture (Rosenblatt and Payton, 1989), and has been used successfully for a number of different simulated mission scenarios including multiple cache retrieval (Huntsberger, 1997), fault tolerance for long duration missions (Huntsberger, 1998), and site preparation (Huntsberger et al., 1999). The system includes all aspects of safety, self-maintenance, and goal achievement that robotic systems require for a sustained planetary surface presence. It currently doesn't include global planning or any adaptive learning capabilities beyond map-making.

The next section briefly describes the organization of BISMARC, followed by a discussion of the action selection mechanism and map-making memory of the system. We close with experimental studies and conclusions.

2. BISMARC Organization

BISMARC is organized as a two level system (shown in Fig. 1). The first level generates possible motor actions using stereo images and the second level uses these action hypotheses coupled with external and internal inputs to drive the actuators on the robot. The *DriveMaps* algorithm used for action generation analyzes local

*The research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contact with the National Aeronautics and Space Administration.

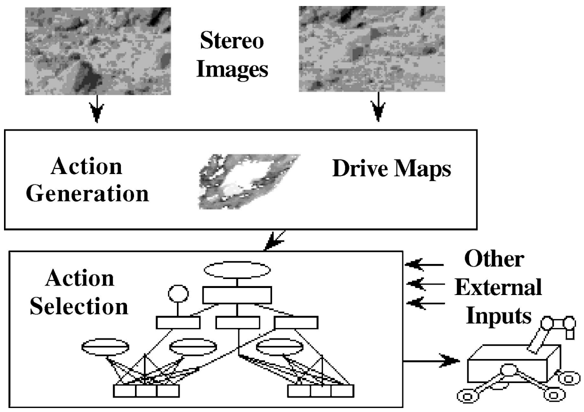

Figure 1. Two level BISMARC architecture with stereo processing action generation, and action selection subsystems.

range information for clear paths relative to a goal and is currently implemented on the SRR and FIDO rovers at JPL. A fuzzy adaptive behavior system with similar capabilities to *DriveMaps* is described in Tunstel (2001).

Figure 2 illustrates an action selection hierarchy for a cache retrieval mission. The rectangular boxes represent behaviors and the ovals are sensory inputs (either fixed, direct, or derived). At the top are the high level behaviors including *Avoid Dangerous Places*, *Sleep at Night*, *Warm Up*, *Scan for Cache*, *Get Cache*, *Cool Down*, *Get Power*, and *Keep Variance Low*. These goals are related to both task and rover safety. For example, since most planetary surface rovers have only visual sensors for navigation, the sensory input for *Proximity to Night* is derived from knowledge of the sun's position and forces the rover to sleep at night by weighting the input to *Sleep at Night* heavier (16.0) than any other behavior in the hierarchy.

The intermediate level behaviors are designed to interact with both the short term memory (STM), which corresponds to perceived sensory stimuli, and the long term memory (LTM), which encodes remembered sensory information. Control loops are prevented through temporal penalties (shown as T-circles in Fig. 2) that constrain the system to only repeat a behavior a predetermined number of times. The bottom level behaviors in the hierarchy fuse the sensory inputs and the activations of the higher level behaviors in order to select an appropriate action to drive the actuators. The next section describes the action selection mechanism of BISMARC in more detail.

3. Action Selection Mechanism

Action selection mechanisms for rovers on a planetary surface require low computational overhead, reactivity even in uncertain environments, no loss of internal state information, combination of conflicting behaviors, and the localization of sensory input to the appropriate modules. The FFH used in BISMARC includes all of these capabilities.

Combining the inputs to a behavioral node is usually calculated as a simple weighted summation. This approach leads to potential problems in the case where the same goal triggers two or more behaviors and the utility of a behavior lower in the hierarchy should not be the sum of their activations. For example, in Fig. 2, the *Get Cache* goal feeds into the *Approach Perceived Cache* and the *Approach Remembered Cache* behaviors. The *Approach Perceived Cache* behavior has a much greater chance of satisfying the goal, since it doesn't involve further travel to a possibly poorly remembered site.

A better activation function, used in BISMARC, balances strong preferences as well as aversion behaviors (i.e., *Avoid Dangerous Places*) (Tyrell, 1993):

$$A_j = \left[\frac{\max_i\left(P_{ij}^{+}\right) + \alpha\square_{i=1}^{N^+} P_{ij}^{+}}{1+\alpha}\right] + \left[\frac{\min_i\left(P_{ij}^{-}\right) + \beta\square_{i=1}^{N^-} P_{ij}^{-}}{1+\beta}\right] \qquad (1)$$

where A_j is the activation output of a node, P_{ij}^{+} and P_{ij}^{-} are the input positive and negative preferences, and α and β are weight factors inversely proportional to the square of N^+ or N^- which are the total number of positive and negative preferences respectively. The activation output of the behaviors associated with directional sensory inputs (shown as segmented ovals in Fig. 2) are multiplied by the sensor stimuli before being used by lower levels in order to suppress activation in directions other than that favored by the sensors.

The weights on the links between modules are heuristically determined based on mission goals. The action selected at the lowest level of the FFH received the highest number of votes. Fuzzy co-norm operators for command fusion at the lowest level of the hierarchy are common to the adaptive behavior system (Tunstel, 2001) and BISMARC. This method is but one of many available (see Pirjanian, 1998 and references therein).

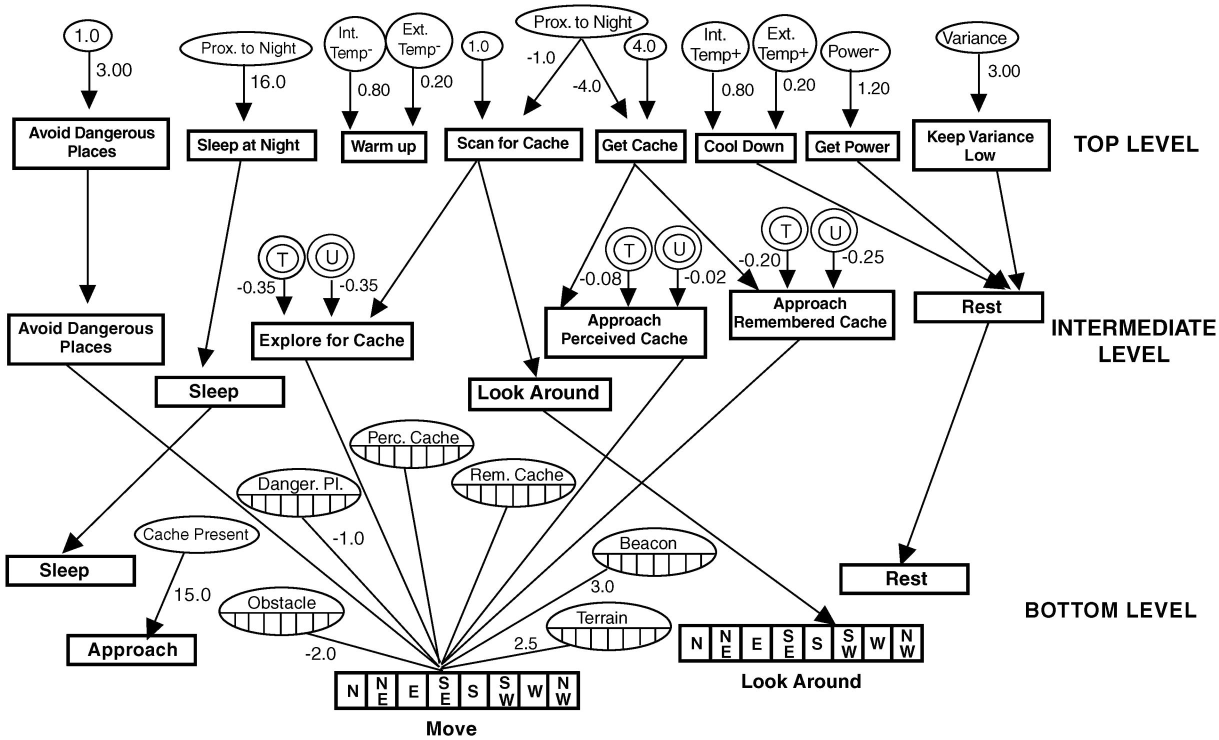

Figure 2. Free-flow hierarchy action selection mechanism for cache retrieval mission scenario. Ovals represent inputs derived from sensory stimuli, rectangular boxex are behaviors, and circles are temporal and uncertainty penalties. All weights on inputs to behaviors are 1.0 unless otherwise noted. Segemented boxes and ovals represent directional inputs (only cardinal directions shown but in practice continuous coverage). See text for further details.

4. Map Making Memory

Work in biologically inspired systems for robotic navigation has a rich history (Mataric, 1991; Kortenkamp and Weymouth, 1994; Touretzky et al., 1994). An overview of other recent studies can be found in Huntsberger and Rose (1998). Biological navigation strategies can be characterized as a four level hierarchy based on complexity analysis (Trullier et al., 1997). The levels are: movement in relation to perception (guidance), orientation with respect to landmarks (place recognition triggered response), movement along known paths (topological), and movement with respect to a global map (metric). The Spatial Semantic Hierarchy (SSH) is a framework that formalizes the relation between the four navigation strategies using five inter-related representations of large-scale space (Kuipers, 2000, 2001). These representations are the sensory, control, causal, topological, and metrical levels. BISMARC embodies aspects of this formalism through the FFH (sensory, control, causal) coupled with the STM (metrical) and LTM (topological) portions of the system.

The STM is distributed in the sense that a detailed occupancy grid (5 cm resolution) (Elfes, 1987) is kept by *DriveMaps* of the last 12.5 m of travel, in addition to the previous and current rover-state vectors (shown as ovals in Fig. 2). These sensor readings are weighted by a perceptual uncertainty based on the absolute difference between the related portions of the previous and current state vectors. A wide swing in values indicates a possibly faulty sensor, which is then monitored and taken off-line if needed.

BISMARC's map-based LTM is similar to that of hippocampus place cells. Landmarks corresponding to obstacles and goals are extensively mapped and stored for comparison to perceived inputs, with a probabilistic update of memories based on the positional variance of the rover and the match strength of the current perception to memory contents. A LTM landmark is encoded as a four-byte field that includes relative height of the landmark (2 bytes), actions leading to the landmark (1 byte), and accelerometer readings on the robot (1 byte). This approach is similar to the coupled goal/representation approach of Mataric (1992) and saves on-board memory use. An alternative approach is an occupancy grid that gives dense coverage of the environment, but doesn't scale well for long duration planetary surface missions.

5. Experimental Studies

In order to determine the utility of BISMARC for complicated planetary surface operations, we have run 2000 simulated multiple cache retrieval missions. A single cache recovery scenario is shown in Fig. 3. The sample acquisition rover filled its cache container until it ran out of power and a second rover with higher speed and mobility recovered the filled container for return to Earth. Our simulation studies used three heterogeneous rovers (single scout and two retrievers) and four cache containers. Mission success was defined as the return of all four cache containers to a landing point. The experimental setup included:

- Random starting and cache positions
- Timestep of 0.1 s
- 10% loss of traction in rocky terrain
- 1 sq. km study area (5 cm resolution)
- Top speed of 15 cm/sec

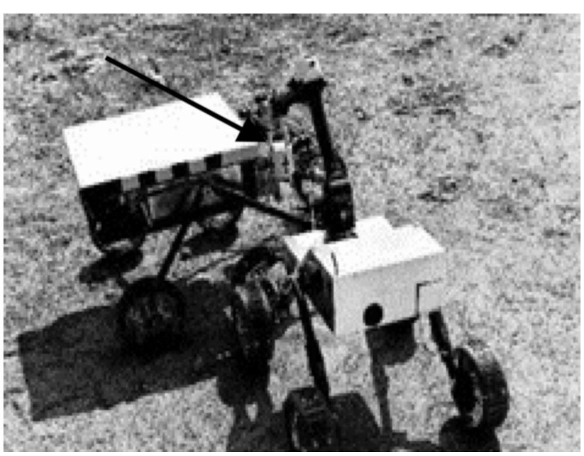

Figure 3. Autonomous rendezvous operation for a cache retrieval scenario in the Arroyo Seco at JPL. Rover on left (Lightweight Survivable Rover) has a cache container (shown with arrow) that is being retrieved using the robotic arm of other rover (Sample Return Rover).

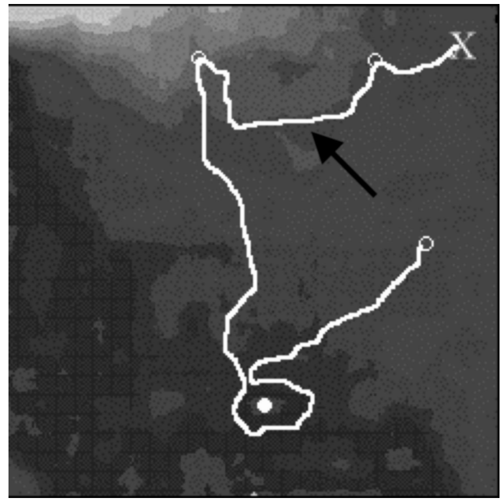

Figure 4. Cache retrieval mission 329, with rover start position at cross and path shown as solid line. Acquired cache containers are shown as open circles and missed container as closed. Study zone is 1 km by 1 km and terrain variation is from −25 m to 300 m.

The scout has three sets of color stereo cameras, a 3 DOF manipulator and a thirty-five week battery lifetime supplemented with solar panels. Each of the two retrievers has two sets of grayscale stereo cameras, a 5 DOF manipulator and a forty week battery lifetime plus solar panels.

A birds-eye-view of mission 329 is shown in Fig. 4, where grayscale is used to indicate height, a cross is the starting position of the scout, the path is a solid line, and the open and filled circles are the cache containers. In this case the temporal penalty prevented the rover from expending its batteries while attempting to satisfy the recovery goal with the cache container in an inaccessible position. Also notable is the portion of the path that skirts the slope (shown by arrow in Fig. 4) on the way to the second cache container. The rover conserved its power by staying on relatively level ground.

Our studies had a 98.9% mission success with no component failures, a 12% success rate with component failures and no fault tolerance, and a 46% success rate with component failures and fault tolerance. An analysis of the 54% of the missions that failed in order to access the impact of a particular component loss gives:

- Loss of tilt sensor(s)—12%
- Loss of stereo sensor(s)—10%
- Loss of wheel encoder(s)—8%
- Loss of battery power indicator—4%

- Loss of internal temperature sensor—2%
- Other (mechanical failure, etc)—18%

This indicates that 28% of the missions failed due to loss of sensors (tilt, stereo, power, internal temperature) that can easily be made redundant within mission constraints. On the other hand, in order to prevent mechanical failures, it is extremely difficult to design redundancy in the mobility sub-systems within mass constraints.

6. Conclusions

We have developed a fault tolerant, autonomous rover control system called BISMARC for multiple planetary rovers. It is based on a blend of local path planning and hierarchical action selection. Fault tolerance was incorporated using a STM update method based on sensor perceptual uncertainty. This method enabled the FFH action selection mechanism to maintain its overall structure even under component failure. Of particular importance for future NASA rover missions was the analysis of the component failures, indicating that an extra 28% of the missions would potentially be successful with single redundancy of these components. Our current directions include automated learning of weights between levels in the action selection hierarchy giving the rovers the ability to adapt to changing environmental conditions such as dust storms, and potential integration into our recently developed CAMPOUT (Control Architecture for Multi-robot Planetary Outposts) running on two prototype rovers at JPL (Pirjanian et al., 2000; Schenker et al., 2000).

Acknowledgments

The author would like to thank the reviewers for their helpful comments and background literature recommendations that greatly improved the completeness and readability of this paper.

References

Agah, A. and Bekey, G.A. 1997. Phylogenetic and ontogenetic learning in a colony of interacting robots. *Autonomous Robots*, 4(1):85–100.

Elfes, A. 1987. Sonar-based real-world mapping and navigation. *IEEE Journal of Robotics and Automation*, RA-3(3):249–265.

Fukuda, T. and Kawauchi, Y. 1993. *Cellular Robotics*, Springer-Verlag: Berlin, pp. 745–782.

Huntsberger, T.L. 1997. Autonomous multi-rover system for complex planetary surface retrieval operations. In *Proc. Sensor Fusion and Decentralized Control in Autonomous Robotic Systems*, SPIE Vol. 3209, pp. 220–229.

Huntsberger, T.L. 1998. Fault-tolerant action selection for planetary rover control. In *Proc. Sensor Fusion and Decentralized Control in Robotic Systems*, SPIE Vol. 3523, pp. 150–156.

Huntsberger, T.L., Mataric, M.J., and Pirjanian, P. 1999. Action selection within the context of a robotic colony. In *Proc. Sensor Fusion and Decentralized Control in Robotic Systems II*, SPIE Vol. 3839, pp. 84–91.

Huntsberger, T.L., Pirjanian, P., and Schenker, P.S. 2001. Robotic outposts as precursors to a manned Mars habitat. In *Proc. Space Technology and Applications International Forum (STAIF-2001)*, pp. 46–51.

Huntsberger, T.L. and Rose, J. 1998. BISMARC. *Neural Networks*, 11(7/8):1497–1510.

Kortenkamp, D. and Weymouth, T. 1994. Topological mapping for mobile robots using a combination of sonar and vision sensing. In *Proc. Twelfth National Conference on Artificial Intelligence (AAAI-94)*.

Kube, C.R. and Zhang, H. 1997. Task modeling in collective robotics. *Autonomous Robots*, 4(1):53–72.

Kuipers, B. 2000. The spatial semantic hierarchy. *Artificial Intelligence*, 119:191–233.

Kuipers, B. 2001. Cognitive maps for planetary rovers. *Autonomous Robots*, this volume.

Maes, P. (Ed.). 1991. *Designing Autonomous Agents Theory and Practice from Biology to Engineering and Back*, MIT Press. Au: Pls. Provide the place of publication

Mataric, M.J. 1991. Navigating with a rat brain: A neurobiologically-inspired model for robot spatial representation. In *From Animals to Animats, Proc. First International Conference on the Simulation of Adaptive Behavior (SAB-90)*, MIT Press, pp. 169–175.

Mataric, M.J. 1992. Integration of representation into goal-driven behavior-based robots. *IEEE Trans. on Robotics and Automation*, 8(3):304–312.

Mataric, M.J. 1997. Behavior-based control: Examples from navigation, learning, and group behavior. *Journal of Experimental and Theoretical Artificial Intelligence, Special Issue on Software Architectures for Physical Agents*, 9(2–3):323–336.

Parker, L.E. 1998. ALLIANCE: An architecture for fault tolerant multi-robot cooperation. *IEEE Trans. on Robotics and Automation*, 14(2).

Pfeifer, R. and Scheier, C. 1999. *Understanding Intelligence*, MIT Press. Au: Pls. Provide the place of publication

Pirjanian, P. 1998. Multiple objective action selection and behavior fusion using voting. Ph.D. Thesis, Laboratory of Image Analysis, Department of Medical Informatics and Image Analysis, Aalborg University, Denmark.

Pirjanian, P., Huntsberger, T.L., Trebi-Ollennu, A., Aghazarian, H., Das, H., Joshi, S., and Schenker, P.S. 2000. CAMPOUT: A control architecture for multi-robot planetary outposts. In *Proc. Symposium on Sensor Fusion and Decentralized Control in Robotic Systems III*, SPIE Vol. 4196, pp. 221–230.

Pirjanian, P. and Mataric, M., 2000. Multiple objective vs. fuzzy behavior coordination. In *Lecture Notes in Computer Science on Fuzzy Logic Techniques for Autonomous Vehicle Navigation*,

D. Drainkov and A. Saffiotti, (Eds.), Springer-Verlag: Berlin, pp. 235–253.

Rosenblatt, J.K. and Payton, D.W. 1989. A fine-grained alternative to the subsumption robot control. In *Proc. IEEE/INNS Joint Conf. on Neural Networks*, pp. 317–324.

Schenker, P.S., Huntsberger, T.L., Pirjanian, P., Trebi-Ollennu, A., Das, H., Joshi, S., Aghazarian, H., Ganino, A.J., Kennedy, B.A., and Garrett, M.S. 2000. Robot work crews for planetary outposts: Close cooperation and coordination of multiple robots. In *Proc. Symposium on Sensor Fusion and Decentralized Control in Robotic Systems III*, SPIE Vol. 4196, pp. 210–220.

Touretzky, D.S., Wan, H.S., and Redish A.D. 1994. Neural representations of space in rats and robots. In *Computational Intelligence: Imitating Life*, J.M. Zurada, R.J. Marks II, and C.J. Robinson (Eds.), IEEE Press, pp. 57–68.

Au: Pls. Provide the place of publication

Trullier, O., Wiener, S., Berthoz, A., and Meyer, J.-A. 1997. Biologically-based artificial navigation systems: Review and prospects. *Progress in Neurobiology*, 51:483–544.

Tunstel, E. 2001. Ethology as an inspiration for adaptive behavior synthesis in autonomous planetary rovers. *Autonomous Robots*, this volume.

Tyrrell, T. 1993. The use of hierarchies for action selection. *Journal of Adaptive Behavior*, 1(4).