PERGAMON

Neural Networks 11 (1998) 1497–1510

Neural Networks

1998 Special Issue

BISMARC: a biologically inspired system for map-based autonomous rover control

Terry Huntsberger*, John Rose

Intelligent Systems Laboratory, Department of Computer Science, University of South Carolina, Columbia, SC 29208, USA

Received 31 October 1997; revised 3 June 1998; accepted 3 June 1998

Abstract

As the complexity of the missions to planetary surfaces increases, so too does the need for autonomous rover systems. This need is complicated by the power, mass and computer storage restrictions on such systems (Miller, *Proceedings SPIE Conference on Cooperative Intelligent Robotics in Space III*, 1829, pp. 472-475, 1992). To address these problems, we have recently developed a system called BISMARC (**B**iologically **I**nspired **S**ystem for **M**ap-based **A**utonomous **R**over **C**ontrol) for planetary missions involving multiple small, lightweight surface rovers (Huntsberger, *Proceedings SPIE Symposium on Sensor Fusion and Decentralized Control in Autonomous Robotic Systems*, pp. 221-227, 1997). BISMARC is capable of cooperative planetary surface retrieval operations such as a multiple cache recovery mission to Mars. The system employs autonomous navigation techniques, behavior-based control for surface retrieval operations, and an action selection mechanism based on a modified form of free flow hierarchy (Rosenblatt and Payton, *Proceedings IEEE/INNS Joint Conference on Neural Networks*, pp. 317-324, 1989). This paper primarily describes the navigation and map-mapping subsystems of BISMARC. They are inspired by some recent studies of London taxi drivers indicating that the right hippocampal region of the brain is activated for path planning but not for landmark identification (Maguire et al., *Journal of Neuroscience*, 17, 7103-7110, 1997). We also report the results of some experimental studies of simulated navigation in planetary environments.

Keywords: Robot navigation; Hippocampal maps; Wavelets; Neural networks

1. Introduction

The recent successful Mars Pathfinder mission has demonstrated that a limited set of geological science tasks can be accomplished remotely using small lightweight rovers. More sophisticated analysis procedures will have to be performed in terrestrial laboratories due to the mass and fragility of the sampling equipment. What this entails is a return of Martian planetary surface samples to the Earth. Such a mission is being planned by NASA in 2005 with the Sample Return Rover (SRR1). SRR1 will retrieve of the cache containers that have been stockpiled by the 2001 and 2003 Long Range Science Rovers (LRSR) during their year-long traversal of the surface, and pass the sample to an orbiter for return to Earth.

Current prototypes of these new rovers, developed at the Jet Propulsion Laboratory, were designed with specific mass and power requirements. The typical mass lies between 7 and 10 kg, and the maximum power use during fast movement (30–50 cm/s) is around 35 W, which can only be sustained for about 6 h without recharging the batteries. The primary sensing modalities on SRR1 include a stereo camera pair (5 cm separation, 130 degree field of view), a goal camera mounted on the manipulator arm (20 degree field of view), an internal gyro and accelerometers, and a planned sun sensor that will give global positioning information. The current SRR1 prototype is shown in Fig. 1, where the manipulator arm is only complete to the first joint and the top of the rover has been removed. Since the available power is shared between computer-related components and drive mechanisms, the active memory capacity is limited to 64 MB of RAM, and a processor with low power requirements such as the PowerPC 603e. Although the hard disk size is currently 500 MB, access is strictly limited to essential needs due to heat dissipation. These limitations constrain the types of algorithms that can be used for sensing and navigation on the planetary surface.

The complicated nature of planetary environments precludes the use of most of the robotic planning systems that are currently available. Although these environments tend to be relatively static, there are often drastic changes in

* Corresponding author. Tel.: +1-803-777-2404; Fax: +1-803-777-3767; E-mail: terry@cs.sc.edu

PII: S0893-6080(98)00088-4

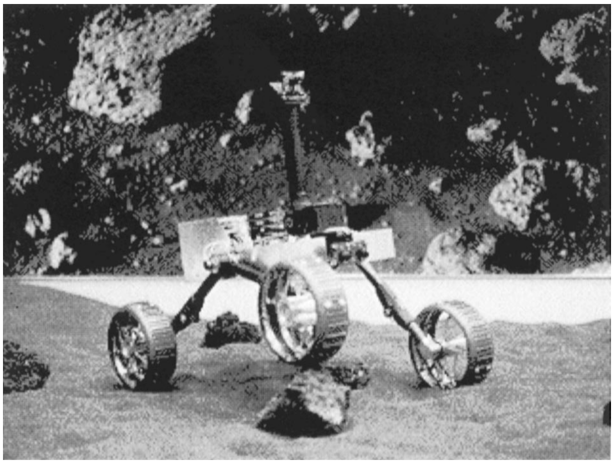

Fig. 1. SRR1 prototype in Planetary Robotics Lab of the NASA Jet Propulsion Laboratory in Pasadena, CA.

altitude, wide temperature variations, and obstacles that are highly irregular in shape. The various internal modules in the rover control system must communicate state information for integrated navigation. Straightforward planning, as opposed to a reactive approach to long-range navigation in such uncertain and harsh environments, would drain the battery reserves within a short time into the mission. The rover prototypes have the option of navigating around an obstacle or riding over it since their ground clearance can be as high as 200 mm. Pure obstacle avoidance would tend to put more stress on the rover structure, since turns involve differential forces. A robust control system needs to be able to combine these possibly conflicting behaviors in order to maximize the mission goal of cache retrieval.

Brooks developed the subsumption control architecture in order to address some of these problems for unstructured terrestrial environments (Brooks, 1986). Using this architecture, seemingly complex behavior arising from a hierarchy of simple, augmented finite state machines is generated. Recent work of Parker has extended this approach and included fault-tolerant characteristics for collections of heterogeneous robots (Parker, 1994). Despite its notable successes, the subsumption architecture loses internal state information due to inhibition or subsumption operations. It also lacks the ability to combine conflicting behaviors due to its winner-take-all strategy for control generation, and is relatively inflexible for highly uncertain environments without reprogramming.

There are also a number of neural network approaches that have been used for robot motion control. Among these are neural dynamics (Baloch and Waxman, 1991; Gaudiano et al., 1997; Hallam et al., 1997), operant conditioning (Bühlmeier and Manteuffel, 1997), reinforcement learning (Kontoravdis et al., 1992; Kröse and van Dam, 1997; Prescott and Mayhew, 1992), backpropagation (Pomerleau, 1991), and self-organization (Heikkonen and Koikkalainen, 1997). Most of these systems are potentially capable of meeting the needs and demands for planetary environments, but have not been tested in that domain.

In an effort to provide a comprehensive set of capabilities suitable for autonomous planetary exploration, this paper presents a multirover control system called BISMARC (**B**iologically **I**nspired **S**ystem for **M**ap-based **A**utonomous **R**over **C**ontrol) (shown in Fig. 2), which is built on previous work that used wolf pack hunting behavior for retrieval of spinning and tumbling satellites (Huntsberger and Hilton, 1995; Huntsberger, 1996). Local rover operations are controlled, using stereo sensor and accelerometer inputs, by a three-level system, which is a hybrid combination of wavelet signal processing, neural networks based on the fuzzy self-organizing feature map (FSOFM) algorithm (Huntsberger and Ajjimarangsee, 1990), and the fine-grained action selection hierarchical network of Rosenblatt and Payton (1989). Our previous work has demonstrated that the FSOFM network is a flexible framework for sensor processing (Huntsberger and Ajjimarangsee, 1990; Huntsberger, 1992a, b, 1995).

BISMARC provides a type of behavior-based control without the need to explicitly program the sensory-to-action behavior mapping as in the original subsumption architecture of Brooks (1986). It also has the ability to adapt to environmental cues because the free-flow hierarchy (FFH) used for action selection is not governed by binary inhibition mechanisms. Experimental studies of 500 missions in a simulated Martian environment using three cooperating rovers demonstrated a 98.9% mission success rate for retrieval of four widely separated cache containers (Huntsberger, 1997). We are not aware of any other study that has simulated missions of similar complexity.

The next section describes biological/biologically inspired models of navigation. This is followed by a discussion of some neural network approaches to robot navigation and control. The overall organization of

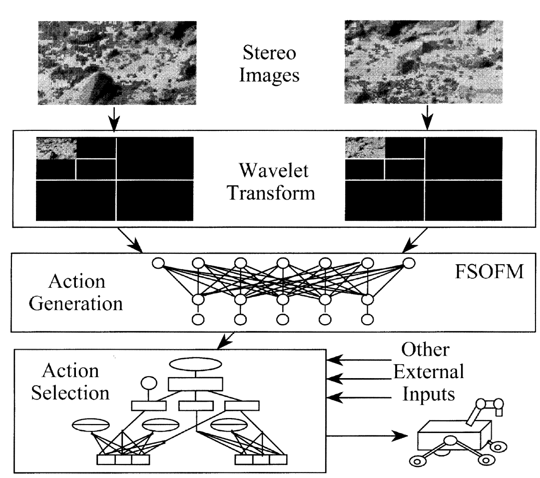

Fig. 2. Multilevel system organization of BISMARC for cache retrieval operations in planetary environments. Coefficients from the wavelet detail channels are used to generate actions with a FSOFM. An ASM then performs a combination operation on the possible actions for final navigation.

BISMARC, including the navigation and map-making portions of the system is detailed in the next section. It should be mentioned here that the *maps* used in our system are sensory-to-action mappings, and that *landmarks* as we use them refer to the actions and internal rover state associated with obstacles, cache containers, and topological features such as cliffs and crevasses. The results of some experiments on simulated multiple cache retrieval missions is described next, followed by a final summary section.

2. Models of navigation

A number of models of navigation are based on theories about the role that the hippocampus plays in spatial map formation and navigation.

Bachelder and Waxman (1994) have developed a biologically motivated, mobile robot visual mapping and localization system. This system uses the Seibert–Waxman 3D object learning and recognition system to learn visual landmarks. Learning a landmark entails the generation of 10–20 aspect categories per landmark. In other words, a landmark is learned by viewing that landmark from 10–20 different directions in 3D space. This allows the system to recognize a landmark when viewing it from different directions. When learning a place, a panoramic view of the environment, which captures information concerning all visible landmarks, is recorded. In this system, places are defined by the locations from which visible aspects of landmarks can be viewed.

Burgess et al. (1994) describe a model of rat hippocampal function in which firing rate maps of hippocampal place cells are treated as approximate radial basis functions over the 2D spatial environment. The phase of firing of individual place cells gives an indication of the relative position of the rat with respect to the spatial location for which the place cell codes. In particular, place cells corresponding to locations in front of the rat fire late in the Theta cycle (the 5–12 Hz EEG Theta rhythm). In contrast, place cells coding for locations through which the rat has traversed fire early in the Theta cycle. The simulation of this model assumes that the search space of 150×150 cm is bounded by four walls. In this model, place cell firing is not strongly modulated by direction. Furthermore, all visual cues are evenly distributed and distinguishable and the issue of movement or removal of individual cues is not addressed.

One view of the function of the hippocampus in navigation is as a path integrator. This view is supported by the ability of rats to navigate in complete darkness suggesting that they are capable of fairly accurate dead reckoning. Samsonovich et al. (1995) suggest a neuron-level model which uses head-direction and ideothetic data to perform path integration which gives the animal's position in space. In this paradigm, landmarks serve the additional purpose of correcting cumulative error in path integration.

The computer model CRAWL developed by Touretzky and Redish (1995) embodies a theory of landmark-based navigation that takes into account behavioral as well as neurophysiological data. Their model combines visual input, head direction, and motor activity to give an estimate of the current location through path integration. Place cells allow a mapping between dead reckoned coordinates and the perceived location of landmarks.

Blum and Abbott (1996) propose a model in which a map of 2D space can be created in the rat hippocampus through a mechanism of long-term potentiation of place cells. The sequential firing of place cells during exploration results in a pattern of long-term potentiation (LTP) in their model which affects subsequent place cell firing. They suggest that the location of the rat is expressed by the ensemble activity of active place cells. Thus, the rat could navigate by moving from its current position to the location corresponding to the strongest place cell activity. They suggest that the coded location would be shifted toward and forward along the specific path that the animal is traversing if LTP occurs while the animal is on that path. The LTP has the effect of evolving the navigational map, extending its range. They evaluate their ideas in a computer simulation of the Morris maze.

In a subsequent paper, Gerstner and Abbott (1997) present a model in which navigational maps are learned through temporally asymmetric potentiation and depression of interacting hippocampal place cell synapses. This model is capable of handling multiple maps to different goals. This is accomplished through the introduction of modulation of receptive fields by goal location. This allows a network of place cells to simultaneously encode maps to several different goals. An interesting aspect of this approach is that modulation corresponding to a goal that has not been previously learned results in a new map that interpolates between learned goal locations to provide a path to the new goal.

A more recent model of map-based navigation proposed by Recce and Harris (1996) applies Marr's theory of hippocampal function. In contrast to many hippocampal-based navigation models, they suggest that many of the additional functional components entailed in map-based navigation are not located in the hippocampus. Specifically, in their model an egocentric spatial map is located in the neocortex and is continuously updated by ideothetic data so that it is possible to have an idea of position in 2D space in the absence of sensory cues. This ability to deduce a homing vector in the absence of specific sensory cues has been observed in animals, but never simulated by models that rely exclusively on these cues. The hippocampus acting as an auto-associative memory stores snapshots of this egocentric map. Head direction cells are used to select the best egocentric map rotation to match the snapshots in the hippocampus which addresses the issue of direction without requiring directional firing patterns in hippocampal place cells. This model is evaluated using a mobile robot in an enclosed internal environment.

Balakrishnan et al. (1997) describe a hippocampal model of spatial learning and navigation for a mobile robot that represents space in a metric framework. In this model, the environment is represented as distinct places in which the center of each place is labeled with metric information derived from path integration. Representing goals using the same coordinate system makes it possible to navigate directly to a goal by finding the vector difference between the current position and the location of the goal. This model combines dead-reckoning position information along with sensory inputs to determine the current location. Thus path integration as well as sensory cues are used to localize the robot in 2D space. They take the approach of using Kalman filter-based tools for analyzing their model for the fusion of uncertain sensory data.

3. Neural network robot control

Traditional approaches to rover navigation follow a sense/plan/act strategy that usually requires long time delays between movements and relatively large computational resources. This technique was successfully applied to cross-country navigation in the CMU Ambler project (Simmons et al., 1991), the JPL Robby project (Gat et al., 1991), and the recent CMU/NASA Ames Nomad trek across the Atacama Desert in Chile (Whittaker et al., 1997). The vehicles used in these projects weighed between 500 and 3000 kg and required up to 2.4 kW of power. The mass and power requirements of long-range rover missions preclude use of such a strategy. These concerns were addressed by subsequent studies undertaken at JPL using reactive methods (Gat et al., 1994; Miller et al., 1992).

The MAVIN system of Baloch and Waxman (1991) as described in the previous section is a comprehensive neural network approach to robotic sensing and navigation. It incorporates an action selection mechanism within the behavior network portion of the system (see Fig. 4 in Baloch and Waxman, 1991). This behavior network is fed by a layer of READ (**RE**current **A**ssociative gated **D**ipole) circuits (see Grossberg and Schmajuk, 1988) that act to associate recognized objects with actions (behaviors), and an emotion network which is based on competitive emotional states. The emotion network is influenced by the READ circuits as well as internal drive inputs such as battery strength and other sensory modalities.

The NETMORC system of Zalama et al. and others (Zalama et al., 1995; Gaudiano et al., 1996a, b, 1997) uses a VAM (**V**ector **A**ssociative **M**ap) (Gaudiano and Grossberg, 1991) to perform unsupervised learning of the mapping of wheel velocities to displacements coupled with the DIRECT (**DI**rection-to-**R**otation **E**ffector **C**ontrol **T**ransform; Bullock et al., 1993) model to learn the mapping of a desired target angle/distance to wheel velocities that provide the necessary movements for goal satisfaction. The advantage of using such a system in uncertain environments lies in its real-time error correction properties. This was exhibited in the experimental studies, where the robot's performance was relatively stable to perturbations of the sensory input, wheel characteristics such as radius and slippage, and noise in the weights connecting the maps (Gaudiano et al., 1996a).

An operant conditioning model was used to extend the NETMORC system to include obstacle avoidance (Gaudiano et al., 1996a). Since the NETMORC system generates the necessary wheel velocities for navigation to a target, the generation of fictitious targets triggered by detection of an obstacle is used for avoidance. The triggering process is modeled based on the operant conditioning model of Grossberg (Grossberg, 1971; Grossberg and Levine, 1987). This model has three cell populations: (1) Short Term Memory (STM) modeled as a recurrent competitive field (Grossberg, 1982) stimulated by sensory inputs (sonar), (2) polyvalent cells driven by inputs from STM and a drive node, and (3) cells connected to a angular velocity map. Robust obstacle avoidance is exhibited by the system due to the learning of the temporal connection between angular velocity patterns and collisions.

The vector field approach used by Tani and Fukumura (1994) requires knowledge of the goal location in order to construct mappings from the sensory inputs into maneuvering outputs. Although training is required for interpretation of the sensory inputs, the use of local minimum in the potential profile derived from the range sensors is particularly effective for obstacle avoidance. A conservative steering equation is used that minimizes the possibility of collisions due to the finite size of the mobile robot. Their experiments demonstrated that knowledge of temporal patterns in the sensory inputs eliminates potential ambiguities in the state space.

The correlation matrix approach of Fukushima et al. (1997) for chained recall of previously learned maps utilizes a shifting stimulus pattern technique based on cross-correlation between the current sensed pattern and the sum of all patterns previously learned in the matrix. Recent studies of London taxi drivers support this framework, where landmark information and spatial maps in path planning activated different areas of the brain (Maguire et al., 1997).

4. BISMARC system organization

The complicated nature of a multiple cache recovery task using multiple rovers involves system organization at many levels. For such operations, we designed the three level BISMARC system (shown in Fig. 2) which uses a hybrid mix of neural networks and behavior-based approaches. The first level performs a wavelet transform on the rover's stereo image pair, the second level inputs these processed images into an action generation navigation network, and then to a third level action selection mechanism (ASM) network modeled after that of Rosenblatt and Payton (1989). Some examples of the other external inputs would include internal temperature sensing, relative time of day, sun sensor positioning, and communications with other rovers.

BISMARC uses a FSOFM to learn *landmarks* (obstacles and goals). In the operational mode, this network generates membership values to the classes of visual input that the system has previously seen. When coupled with onboard rover components such as accelerometers and dead reckoning inputs, an egocentric *map* of the environment is built using the FSOFM response as an index. This would be equivalent to navigation based on a statement such as 'approximately five steps past the big oak tree', which does not indicate which big oak tree, but would allow localization after an oak tree is sensed. In the same manner, a human would have trouble distinguishing between navigation in a real environment and in an accurate Hollywood movie set of the same environment. Since unsupervised operation or training a system in a planetary environment such as Mars would be costly and potentially dangerous to the rover, BISMARC offers a compromise solution.

The three-level self-organizing system of Heikkonen and Koikkalainen (1997) is closest to BISMARC in structure. They use Gabor filtered visual input fed into the first level, followed by clustered feature extraction in the second level, and action generation in the third level. Lack of a specific action selection mechanism in their system is the most notable difference between the two. Our system replaces the first two levels of their system with a single FSOFM by using the wavelet transform as a preprocessing level. This was done in order to eliminate the large storage space needs and coverage problems of the frequency space by the Gabor wavelet (Lee, 1996). We use the FFH of Rosenblatt and Payton (1989) for action selection, which had been successfully fielded for an autonomous land vehicle (ALV). The FFH used in our study includes both unidirectional and multidirectional sensors, temporal penalties for actions (Sutton, 1988), ordinary and multidirectional nodes, and both additive and multiplicative processing of inputs to a node.

This type of FFH was recently shown to be optimal within the multiple objective decision-making (MODM) formalism, which produces an action through the maximization of a global objective function that includes all possible actions (Pirjanian and Christensen, 1997). Additive processing of inputs to a node can combine conflicting behaviors for maximum motion efficiency (shown in Fig. 3). The rover would move south if the *Avoid Obstacle* action generation were the only node that is activated since it has a high negative strength as represented by the large shaded circle in the north direction. Instead its combination with the *Approach Goal* activation, which contains a small positive value, results in motion towards the cache container perceived to the west.

Although the rover(s) that place the caches have some sense of position, they will not generally have access to a detailed map. This would be the case even with known landmarks due to uncertainty in accuracy of localization using triangulation (Sutherland and Thompson, 1994). Field trials with the current generation of Lightweight Survivable Rovers (LSR) and SRR1 at the Jet Propulsion Laboratory have indicated that there is a great amount of slippage in the drive wheels during traversal of Martian-like terrain. This precludes the use of purely dead-reckoning techniques for navigation and obstacle avoidance. This point was recently addressed by Baumgartner and Skaar using an extended Kalman filter analysis of staged visual cues in the environment to refine the dead reckoning position estimates (Baumgartner and Skaar, 1994). Their system accurately localized a rover in the laboratory to within 1 cm using only a single camera and a series of previously defined reference patterns. Additional work on this problem by Roumeliotis and Bekey indicates that the extended Kalman filter analysis can be further improved by an order of

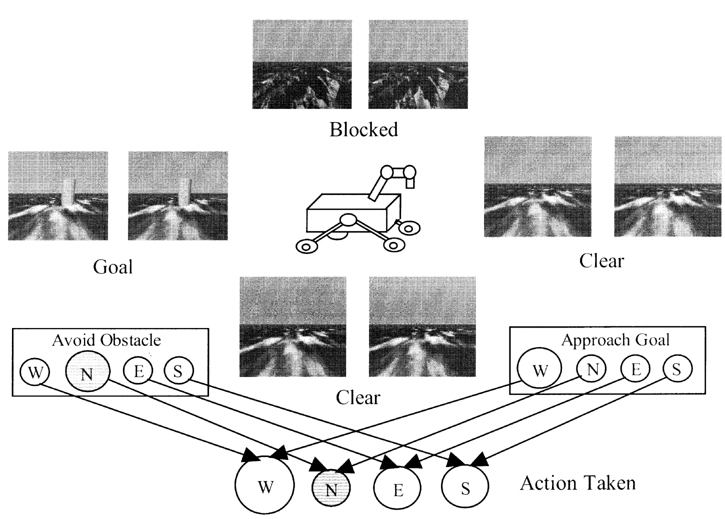

Fig. 3. Combination of conflicting behaviors where pure obstacle avoidance action generation leads to inefficient motion away from the cache container. Shaded circle indicates a negative value and relative size of circles indicate strength.

magnitude with the inclusion of a sun sensor input (Roumeliotis and Bekey, 1997). The ASM used in BISMARC combines the information from the sun sensor and the stereo cameras on the scout rover to extract visual cues from the environment, coupled with a FSOFM to perform the kinematic adjustments. The output of the FSOFM serves as a behavior based index into a sensory-to-action map as shown in Fig. 4, where a relatively high level of sensory-to-action detail is only maintained at obstacles and goals.

The rovers in the simulated missions were of two types: scout and retrieval. The navigation system on the scout rover is responsible for locating the cache containers and 'marking' obstacles that it encounters in its path to each of the containers using the sensory-to-action information generated by the ASM level in BISMARC. These obstacles need to be avoided both by the scout rover, and the retrieval rovers on their way to each cache with their augmented knowledge of the terrain. This process requires an effective representation of the terrain, including obstacles and the cache container locations. The use of stereo cameras for characterization of the terrain offers the possibility of real-time rover obstacle avoidance (Matthies, 1992).

Many biological systems use anchor points (landmarks) for navigation in spatially extended environments. This was recently demonstrated in a study of London taxi drivers where the right hippocampus was activated for route planning, but not for landmark identification (Maguire et al., 1997). The medial parietal cortex was activated for both tasks, indicating that landmark identification is an important portion of the egocentric sequential aspects of the route planning task (Maguire et al., 1997). In other words, navigation between two waypoints that are not in a direct line involves the sequential reference to landmarks along the way. This is the primary motivation for the sensory-to-action maps that are built for navigation in BISMARC.

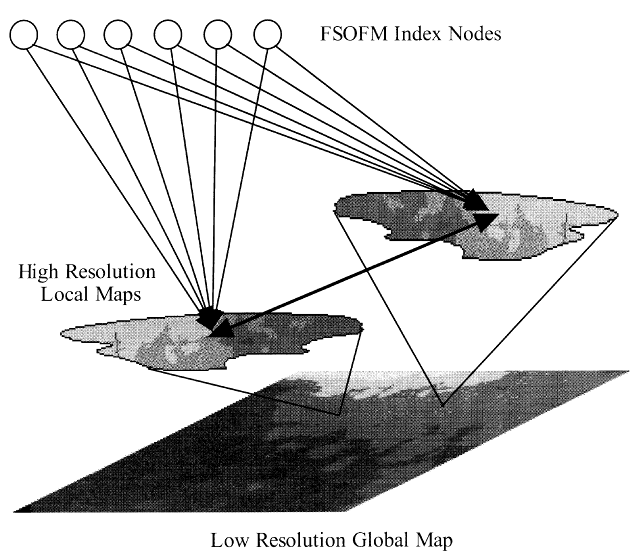

Fig. 4. Sensory-to-action index into local maps. These local maps are only obtained at obstacles and goals and are in an order of magnitude better spatial resolution than the global map.

We have developed a landmark identification subsystem in BISMARC based on the FSOFM, which generates behaviors associated with wavelet-encoded visual input. It should be stressed here once again that *landmarks* in BISMARC are an encoded action representation, and are not meant in the traditional sense (i.e. the Statue of Liberty). This subsystem is similar to the approach taken by Matarić (1992a, b, 1997) except that her studies used a behavior chain to characterize a path through the indoor environment. The output of the FSOFM is used as a local index into a crude map of the terrain (5 m resolution) that would be obtained with a camera as the rover landed. The indexing process is shown in Fig. 4.

Navigation with such a low resolution map would be extremely dangerous for the rover, since the rover size is on the order of 50 cm. BISMARC uses a 5 cm resolution for the local maps, which would correspond to a frequency of 10 samples/s when moving at full speed. The crude map used in BISMARC is a metric map of the terrain with a 4 bit label at each grid point indicating the possible terrain character (texture measure of roughness). Due to the unique landing mechanism used in the Sojourner mission, acquisition of this type of visual 2D map was not possible. Future NASA mission plans include a more conventional landing scenario. Links between the landmark points are encoded using the odometry data, which is used for rough dead reckoning analysis of the connecting path. Although wheel slippage and other effects make these data unreliable, significant information compression is achieved by not storing details of the uninteresting (at least to these rovers) portions of the path.

Since the area that is traversed is on the order of square kilometers, a detailed occupancy grid approach like that of Moravec and Elfes (Elfes, 1987; Moravec and Elfes, 1985) is not feasible due to storage limitations on the rovers. For example, our simulation studies covered a 1 km $\times$ 1 km area at a local resolution of 5 cm. This would require a grid of 20,000 $\times$ 20,000 places for full coverage, or about 400 MB if a single byte is used for each grid point. Extensive access of the hard disk would be required for such a map to be kept, which would lead to possible internal temperature problems for the rover. BISMARC only encodes the local details of the surface that are relevant for navigation, which include the relative height of the obstacles, action taken immediately prior to and after encountering the obstacles (direction of travel and velocity), accelerometer state (to encode surface slope), and a detailed 360 degree sweep of the area at any cache container. This representation uses four bytes for obstacles and twelve bytes for the cache. This information is only stored for obstacles and goals (cache containers).

A vision preprocessing level uses wavelet-based algorithms to decompose the images generated by each of the stereo cameras. The wavelet decomposition provides information about the scale, location, and orientation of features in an image. Because the wavelet decomposition contains information about the local frequency content of an image, it can represent visually important features (such as edges) more compactly than many of the other transforms commonly used in image processing (Mallat and Zhong, 1992). Our previous work with wavelet-processed images has shown their utility for sensor fusion (Huntsberger and Jawerth, 1993), morphological image processing (Huntsberger and Jawerth, 1995), motion analysis (Huntsberger et al., 1994), face processing (Huntsberger et al., 1998), image enhancement (Hilton et al., 1993), and texture processing (Espinal et al., 1998).

After the wavelet transform is performed on each of the stereo images, a vector is formed using the multiresolution information from the two highest levels of the wavelet horizontal and vertical detail channels. This vector of length 40,960 elements is the input to the FSOFM, the output being any of six action states: go forward, backup, turn right, turn left, stop, or pick either direction in turn. This sensory-to-action mapping approach is similar to that used by Pomerleau in the CMU ALVINN road following system (Pomerleau, 1991), although in that system specific features such as edges of roads were identified. BISMARC is only encoding the raw stereo visual information without any attempt to label individual features or objects beyond the desired action associated with the input pattern. Although the FSOFM is an unsupervised network, it can be trained by presenting it with samples and labeling the output nodes that correspond to each of the generated actions. The weights are then clamped for all subsequent runs. This approach was successfully used for the inverse kinematics study of a seven degree-of-freedom robot arm (Soltys and Huntsberger, 1993).

The network is trained with a set of wavelet-processed images that simulate the types of terrain seen on the Martian surface during the recent Sojourner mission. Some of the stereo pairs that were used to train the navigation FSOFM are shown in Fig. 5. A total of 500 stereo pairs were used for the training session, which took 637 epochs to converge. This should be compared to a backpropagation implementation, which typically takes hundreds of thousands of epochs to converge. Recall of the trained images was 100%.

An advantage of using the FSOFM for the action generation level lies in the membership values that are generated at the output nodes. The sum of these values is normalized to one, and the relative size of the membership values gives a ranking of the actions that are possible. For example, a set of stereo images shown in Fig. 6 that were not in the training set were shown to the system after training. This set of images contains a cache container that is partially hidden by rocks on the right-hand side of the field of view. The

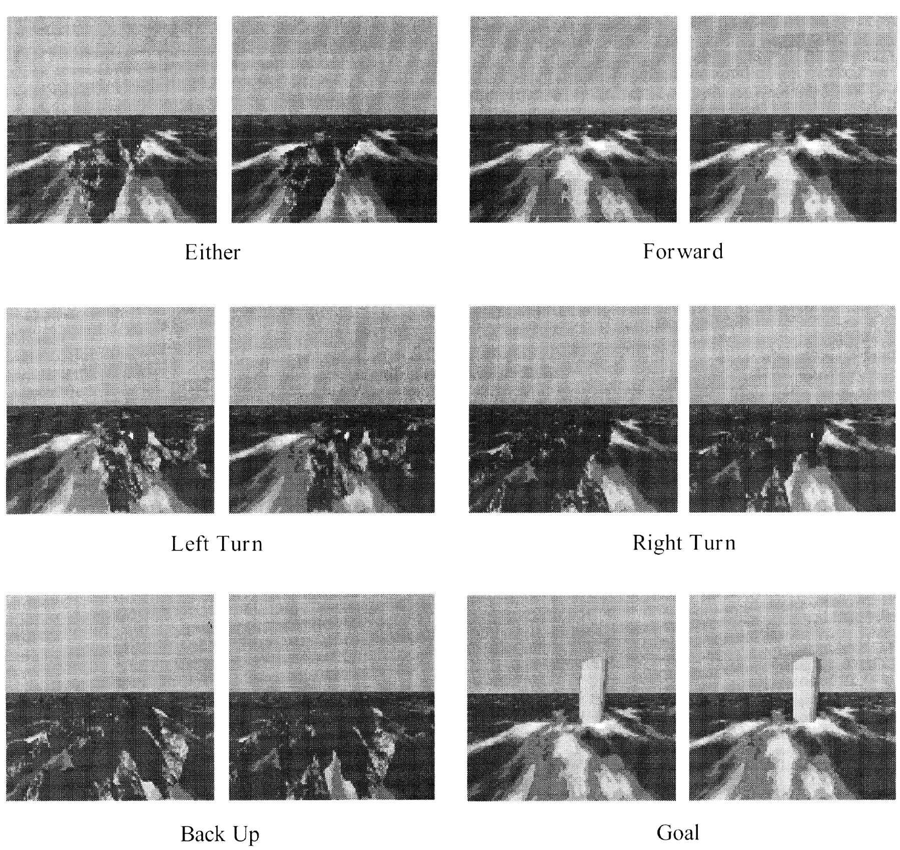

Fig. 5. Samples of stereo image pairs used to train the navigation FSOFM; generated action is given beneath each pair. Sample stereo pairs generated using VistaPro with USGS Martian Digital Elevation Maps.

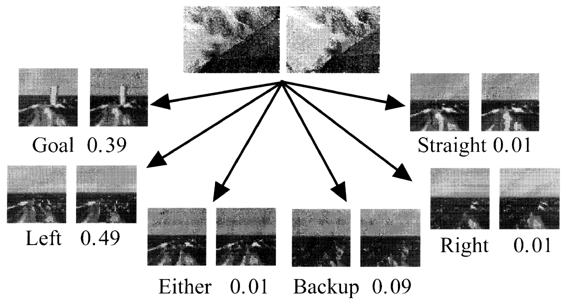

Fig. 6. Membership values returned from FSOFM with ambiguous input. Although the system was not trained with this stereo pair, it was able to generalize to the appropriate action of turn left to avoid the obstacle.

ranking of the membership values indicated the highest actions were to turn left (0.49), goal (0.39), with the rest of the memberships distributed over the other four states. The obstacle avoidance behavior of the system is invoked due to it having the highest membership value, but the system still senses that the goal is close as it moves. Movement far from an obstacle is automatically generated, and the rover will only turn when its field of view has the vector of wavelet coefficients that indicate looming of an obstacle large enough to need avoidance behavior. The velocity of the rover is selected within the ASM dependent on the FSOFM output, it is slower during turns (3 cm/s), when approaching the goal (1 cm/s), and in rocky terrain (2 cm/s).

There are a number of factors beyond the visual sensory input that influence navigation of the rover. These include the health of the rover (internal/external temperature, battery power levels, accelerometers/gyro), time of day, and homogeneity of the terrain. For example, if the rover is approaching a highly uneven portion of the terrain rather late in the day, the control decision may be made to halt and wait for the next morning in order to recharge the batteries and to have enough light for visual sensing. Action selection in the presence of such conflicting behaviors is done in BISMARC using the FFH of Rosenblatt and Payton (1989).

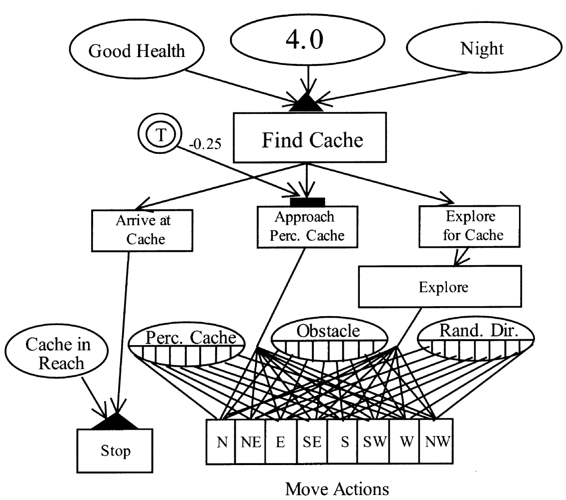

Fig. 7. *Find Cache* system for BISMARC. Connections from other subsystems are not shown. Notation for symbols is that of Tyrrell (1993). See text for a detailed discussion.

Possible evidence for a type of FFH behavior in nature is found in geese when offered corn from a human hand. Fear of the human coupled with the desire for the corn leads to a quivering of the neck due to antagonistic muscle responses (Lorenz, 1981). A FFH is a directed graph of action and stimulus nodes that are combined using predetermined rules. These rules may include addition, multiplication, or more complicated means of combination. The *Find Cache* system for BISMARC is shown in Fig. 7. Action nodes are drawn as rectangles, stimulus nodes as ellipses, and those with multidirectional characteristics are indicated using 8 directional bins. The combination rules are additive for a small filled rectangle above the node, multiplicative for a small filled triangle, and the following rule is used for plain rectangular nodes (Tyrrell, 1993):

$$A_j = \frac{\max_i(P_{ij}^+) + \alpha \sum_{i=1}^{N^+} (P_{ij}^+)}{1+\alpha} + \frac{\min_i(P_{ij}^-) + \beta \sum_{i=1}^{N^-} (P_{ij}^-)}{1+\beta},$$

where A_j is the activation strength, P_{ij}^+ and P_{ij}^- are the positive and negative preferences from node i for node j, N^+ and N^- are the numbers of such preferences for node j, and α and β are constants. This more sophisticated combination rule was developed by Tyrrell (1993) to guarantee the proper transfer of goal and motivational behavior to lower levels of the FFH.

Since the scout rovers in our simulations need to find all of the cache containers, penalty nodes were added to prevent a total mission failure in the event that one or more of the containers are impossible to approach. This scenario would have occurred during our simulation runs (Mission 329 shown in Fig. 10), leading to a circling behavior around the cache container until the rover's batteries were exhausted. Tyrrell (1993) introduced the temporal penalty (T-circle in Fig. 7) to inhibit action that will take an inordinate amount of time to complete. The temporal penalty is derived using the assigned value raised to the power of the elapsed time during the current action. Since it is being used with a multiplicative combination rule and is less than 1.0 to start, the action activation will decline with increasing time. This type of node is similar to the *impatience* parameter used by Parker (1994) or the *temporal discount* factors of Sutton (1988). Temporal penalty nodes increased the likelihood of satisfying the overall mission goal of maximizing the number of cache containers returned.

The large motivational input of 4.0 to the *Find Cache* node is mediated through its multiplicative combination with the internal stimuli of *Night* and *Good Health*. The *Night* stimuli is derived from the sun sensor input, and is equal to 1.0 at sunrise and decreases to 0.0 at sunset. The *Good Health* stimuli is a product of the battery level scaled between 0.0 and 1.0 and a triangular function of the internal temperature of the rover. This triangular function starts at

0.0 for −50 °C, peaks at 1.0 for 55 °C, and drops to 0.0 at 100 °C. The multidirectional stimuli that feed the *Move Actions* output node are directly derived from the output nodes of the FSOFM.

While the design of BISMARC has been biologically inspired, the constraints placed on its design, which are a consequence of the environment in which it must operate, set it apart from the models described in the previous two sections. Any autonomous navigation model for a planetary rover is a priori expected to operate in the open in 3D space. Furthermore, such a remote system does not have the luxury of being able to learn to navigate or learn a map of the environment in situ by bumping into obstacles or careening over a precipice. The rover must avoid self-inflicted damage at all cost due to the paucity of rover repair facilities and qualified technicians on remote planets. Prior knowledge of the exact goal position in the environment will not be known due to the uncertainties in navigation mentioned in the previous section. Additionally, time, computation and power constraints also place limits on map learning.

The sensory input to BISMARC is also quite different from the systems that were detailed in the preceding sections. Since it operates in 3D space, it also uses accelerometer data to indicate whether its vertical orientation is level, downslope or upslope. In addition to ordinary visual information, the visual information in its field of view is qualitatively characterized as left-field occluded, right-field occluded, total occlusion, goal, clear, and centered.

The direct integration of the obstacle avoidance behavior into BISMARC's route planning strategy is a necessary component of the design. Since obstacles are used as *landmarks* for sensory/action map making, detailed information about size, relative height, etc., are important for the retrieval rovers that will follow the scout. The necessary clearance for a rover to navigate around an obstacle is built into the training sets, and looming is used to differentiate between small obstacles that can be driven over and the larger ones that require course changes. Most of the systems reviewed below in the context of obstacle avoidance are different from BISMARC in the sense that they are actively mapping the environment using obstacles as something to avoid, rather than as significant features for landmark-based navigation.

The MAVIN system (Bachelder and Waxman, 1994) uses 2D information for landmark recognition and has been operated in an enclosed environment. However, there are no constraints imposed on the MAVIN system which preclude its operation in an open 2D environment.

The 2D models of Burgess et al. (1994), Samsonovich et al. (1995), Blum and Abbott (1996), Gerstner and Abbott (1997), and Balakrishnan et al. (1997) have been simulated in enclosed environments. The lack of published data showing the performance of these models in controlled open environments makes it difficult to draw objective conclusions regarding potential performance in arbitrary open environments. This is particularly difficult to evaluate in the cases of the 2D models of Burgess et al. (1994), Blum and Abbott (1996), Gerstner and Abbott (1997) in which explicit use is made of the boundaries of the closed environment in the simulations. The approach of allowing the simulated rat to 'bounce' off of walls has the effect of redirecting the focus of exploration toward the interior of the enclosed environment. Nevertheless, the fact that these systems, as well as those by Touretzky and Redish (1995) and Recce and Harris (1996), have only been tested or simulated in enclosed environments should not be interpreted to mean that such systems are only capable of operating in an enclosed environment.

Burgess et al. (1994) state that a consequence of the sensory input entorhinal cell model that they employ is that the goal location must be inside a convex hull of the set of cues in order for the navigation to be successful. This could pose a problem in an arbitrary open environment.

In the case of the 2D model of Samsonovich et al. (1995), the connections between place cells in the model are prewired and fixed. In addition, the description of the model does not address obstacle avoidance (Samsonovich et al., 1995). At a minimum, navigation in an open/unbounded environment would require the ability to avoid obstacles unless the map is prelearned and the environment is static.

The 2D models of Burgess et al. (1994), Blum and Abbott (1996), Gerstner and Abbott (1997) address obstacle avoidance by allowing the simulated rat to 'bounce' off of the enclosing walls or any obstacles. While such 'bouncing' may not be unreasonable in the simulation of a model, it is clearly unacceptable in the case of a navigational system for an actual rover. Since our goal is the latter, we have been forced to take a different tack in addressing obstacle avoidance.

The CRAWL model by Touretzky and Redish (1995) is also a 2D space simulation, although the authors indicate that portions of the model have been implemented on a mobile robot in an enclosed environment. They describe computer simulations of a split landmark array task. Their model is trained on the environment by depositing the simulated animal at successive random spots. Place units are recruited as necessary to learn the environment with respect to randomly chosen visible landmarks. Training terminates after 10 random spots have been visited without requiring the recruitment of additional place units. A consequence of this training procedure is that the issue of obstacle avoidance does not arise. Although they do not present results for an open environment, there does not appear to be anything in their model that would preclude such experiments assuming that something other than random site selection were used to train the model.

Recce and Harris (1996) evaluated their model using an actual mobile robot, ARNE in an enclosed 2D environment. Irregular shaped rooms were used in their experiments to reduce the similarity of features that would be observed with the robot's sonar sensor. However, the rooms were large enough so that the robot was not able to detect all walls from one location. The exploration strategy that is described is a simple wall-following strategy. Thus the boundaries of the enclosed environment play an important role in the results that are presented. However, there does not appear to be any intrinsic limitation in their architecture which would preclude the extension of their model to address an open environment given an exploration strategy not predicated on an enclosed environment.

5. Experimental study

The interface to BISMARC (shown in Fig. 8) has display output for relevant mission parameters, including the current internal temperature and battery level for the rover, a bird's-eye-view of the study area, the current stereo view from the rover, the current top two actions that have been activated, the Martian time of day, and the total accrued mission time. The path that has been followed by the rover up to this point in the study mission is shown as a white line on the bird's-eye-view display.

The mission control strategy is as follows:

- Deployment of all rovers at lander staging site.
- Scout rover heads toward first cache site, sending back coarse path information such as major obstacles and path lengths between them. It is followed by a single retrieval rover using the *follow* behavior of Matarić (1992a, b). In addition wavelet-compressed images of selected views to augment the map and of the cache site are broadcast to the retrieval rovers when the scout rover arrives. This is done to safeguard the mission objectives in the event that the retrieval rover does not make it to the cache site.
- After deployment of a beacon, the scout rover starts toward the next potential cache site location, once again broadcasting compressed versions of its path. If another retrieval rover is available, it is dispatched on an intersection path with the scout.
- Each retrieval rover broadcasts a message indicating successful acquisition of the cache and another when it returns to the lander staging site. The *robot impatience* behavior (Parker, 1994) will be invoked if the waiting retrieval rovers do not receive this message within a predetermined period of time.
- After all known caches are retrieved, the rovers attempt to return to the lander staging site.

This control strategy was coded as a FFH, and run under BISMARC.

We ran 500 trials using a subset of the heightfield information returned from the Mars Pathfinder mission (shown in Fig. 9). The area encompassed about 1 km × 1 km with a grid decomposition resolution of 5 cm at the detailed map level. The number of detailed maps will vary with the ruggedness of the terrain, since the maps are only generated for obstacles and cache containers. Each trial had different starting positions and the placement of 4 cache containers was randomized within the area. It was assumed that each cache placement site was known to within a 200 m radius. Three rovers were deployed for each simulated mission: a scout and two retrieval rovers. The bandwidth of the

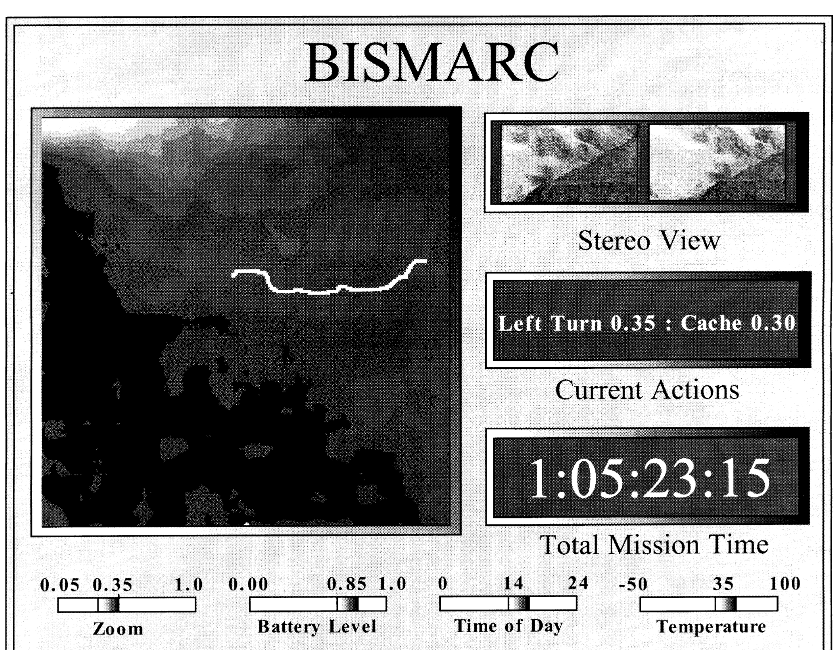

Fig. 8. Interface to BISMARC. Bird's-eye-view of study area displayed in left subwindow (grayscale is used to indicate relative elevation), stereo view from rover displayed in upper right subwindow. Important mission parameters such as internal temperature and battery level shown as dynamic sliders.

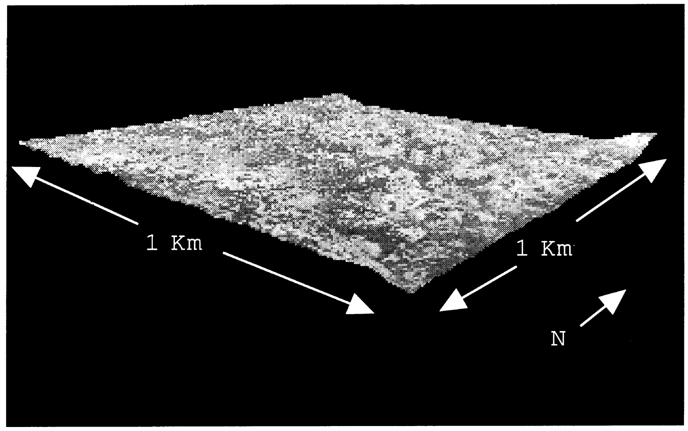

Fig. 9. 1 km × 1 km study zone used in the experimental studies. Grayscale is used to indicate relative elevation, with a highest elevation change in the study area of 500 m.

communication channel between the rovers is 1 Mbit/s, which is the same as the modem installed in the current SRR1 prototype at the Jet Propulsion Laboratory. The top speed on the rovers was set at 30 cm/s, which is consistent with the SRR1. In order to simulate wheel slippage, we set a 5% loss of traction when climbing over a rock or traversing rocky terrain. The battery lifetime was set at one week on all of the rovers and the timestep size for the simulations was fixed at 0.1 s. All of the rovers were forced to sleep during the night hours of the simulations, since there were no infrared sensors on any of the rovers, and navigation at night would be dangerous.

A total mission success is defined as the return of all four of the cache containers to the landing site. The mission success rate for the 500 simulated missions was 98.9%. Only one of the missions was a total failure, in that none

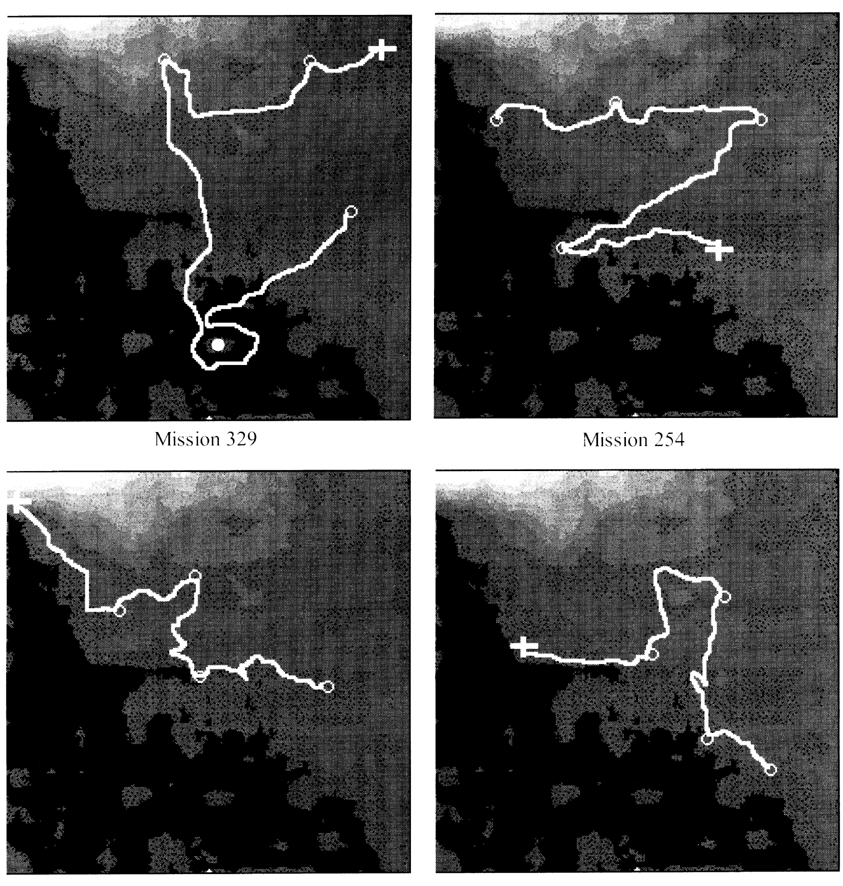

Fig. 10. Paths taken by the scout rovers in four of the multiple cache retrieval missions. Cache container positions are marked by circles and rover start position is marked with a cross. Filled circle in Mission 329 is a missed cache container. Grayscale is used to indicate relative elevation.

of the cache containers were recovered due to the untimely death of the scout rover from a fall down a steep incline. The paths taken for four of the missions are shown in Fig. 10, where the white line in each subwindow is the path and the cache container positions are marked with white circles. An example of the temporal penalty effect on rover navigation is seen in Mission 329 in Fig. 10, where one of the cache containers (shown as a filled circle) was unable to be acquired due to it being totally surrounded by rocky terrain. It is unclear whether this scenario would ever occur in an actual mission, since the cache container was supposedly put in place by the 2001 or 2003 rovers. The question raised is how those rovers got into this area of the planetary surface in the first place. Evaluation of the optimal path for each mission is complex due to the obstacle avoidance options in such a complicated 3D terrain, and only a complete path planning analysis would give this information. In general, the scout rovers localized the cache containers well within the mission time constraints and with a minimal amount of structural damage. Most of the additional path length can be traced to the *Explore* behavior in the *Find Cache* ASM of Fig. 7, which was motivated using random search directions. Current NASA plans call for a beacon guidance system for the initial 200 m approach phase, followed by visual guidance to the cache container when within 10 m.

6. Conclusions

This paper introduced a hybrid system for planetary rover navigation and control called BISMARC that uses neural networks coupled with behavior-based control for the coordination of multiple rovers during the navigation to and subsequent retrieval of multiple cache containers. Out of a total of 2000 cache containers, 1978 were successfully retrieved over 500 simulated missions. These results indicate that the sensory-to-action map representation used by BISMARC and shared between the rovers was an effective representation of the study area for retrieval operations. The majority of the failures were due to expiration of the battery lifetime on the rovers. The average completion time over the 500 missions was 4.7 days.

We are currently extending the navigation algorithms to include better search strategies in order to minimize the chance of loss of power leading to mission failure. An investigation is also underway to determine the length of the optimal vector of wavelet coefficients in order to reduce the FSOFM processing time. It is anticipated that field studies at the Jet Propulsion Laboratory will be used to test the navigation portion of the system for SRR1 in real terrain. In addition, we are evaluating incorporation of more fault tolerance into the third level network of BISMARC to maintain a reduced mission capability in the event of internal sensor failure or partial structural damage to the rover.

Acknowledgements

The authors would like to thank the reviewers for their extremely thorough reading of the manuscript and suggestions which greatly improved its readability. The first author would like to thank Dr. Paul Schenker of the Jet Propulsion Laboratory for his support under the ASEE NASA Summer Faculty Fellowship Program during the summer of 1997. This work was also supported in part under ARO grant DAAH03-96-C-0034, and ONR grant N00014-97-1-1163.

References

Bachelder I.A., & Waxman A.M. (1994). Mobile robot visual mapping and localization: a view-based neurocomputational architecture that emulates hippocampal place learning. *Neural Networks*, *7 (6) 7*, 1083–1100.

Balakrishnan, K., Bousquet, O., & Honavar, V. (1997). Spatial learning and localization in animals: a computational model and its implications for mobile robots. Technical Report 9720. Ames, IA: Iowa State University.

Baloch A.A., & Waxman A.M. (1991). Visual learning, adaptive expectations, and behavioral conditioning of the mobile robot MAVIN. *Neural Networks*, *4*, 271–302.

Baumgartner E.T., & Skaar S.B. (1994). An autonomous vision-based mobile robot. *IEEE Transactions on Automatic Control*, *39 (4)*, 493–502.

Blum K.I., & Abbott L.F. (1996). A model of spatial map formation in the hippocampus of the rat. *Neural Comp.*, *8*, 85–93.

Brooks R.A. (1986). A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, *2 (1)*, 14–23.

Bühlmeier, A., & Manteuffel, G. (1997). Operant conditioning in robots. In O. Omidvar & P. van der Smagt (Eds.), *Neural systems for robotics*. San Diego, CA: Academic Press.

Bullock D., Grossberg S., & Guenther F. (1993). A self-organizing neural network model for redundant sensory–motor control, motor equivalence, and tool use. *Journal of Cognitive Neuroscience*, *5 (4)*, 408–435.

Burgess N., Recce M., & O'Keefe J. (1994). A model of hippocampal function. *Neural Networks*, *7 (6) 7*, 1065–1082.

Elfes A. (1987). Sonar-based real-world mapping and navigation. *IEEE Journal of Robotics and Automation*, *RA3 (3)*, 249–265.

Espinal F., Huntsberger T., Jawerth B., & Kubota T. (1998). Wavelet-based fractal signature analysis for automatic target recognition. *Optical Engineering*, *37 (1)*, 166–174.

Fukushima K., Yamaguchi Y., & Okada M. (1997). Neural network models of spatial memory: associative recall of maps. *Neural Networks*, *10 (7)*, 971–980.

Gat E., Desai R., Ivlev R., Loch J., & Miller D.P. (1994). Behavior control for robotic exploration of planetary surfaces. *IEEE Journal of Robotics and Automation*, *10 (4)*, 490–503.

Gat, E., Slack, M., Miller, D., & Firby, R. (1991). Path planning and execution monitoring for a planetary rover. In Iyengar & Elfes (Eds.), *Autonomous mobile robots: Control planning and architectures*. IEEE Computer Society Press.

Gaudiano P., & Grossberg S. (1991). Vector associative maps: unsupervised real-time error-based learning and control of movement trajectories. *Neural Networks*, *4 (2)*, 147–183.

Gaudiano, P., Guenther, F.H., & Zalama, E. (1997). The neural dynamics approach to sensory–motor control. In O. Omidvar & P. van der Smagt (Eds.), *Neural systems for robotics*. San Diego, CA: Academic Press.

Gaudiano, P., Zalama, E., Chang, C., & Coronado, J.L. (1996a). A model of operant conditioning for adaptive obstacle avoidance. Technical Report. Boston, MA: Boston University.

Gaudiano, P., Zalama, E., & Coronado, J.L. (1996b). An unsupervised neural network for low-level control of a wheeled mobile robot: noise resistance, stability, and hardware implementation. *IEEE Trans. Systems, Man, Cybernetics., 26,* 485–496.

Gerstner W., & Abbott L.F. (1997). Learning navigational maps through potentiation and modulation of hippocampal place cells. *J. Computational Neurosci.*, *4*, 79–94.

Grossberg S. (1971). On the dynamics of operant conditioning. *Journal of Theoretical Biology*, *33*, 225–255.

Grossberg, S. (1982). In S. Grossberg (Ed.), *Studies of mind and brain: Neural principles of learning, perception, development, cognition, and motor control*. Boston, MA: Reidel.

Grossberg S., & Levine D. (1987). Neural dynamics of attentionally modulated Pavlovian conditioning: blocking, interstimulus interval, and secondary reinforcement. *Applied Optics*, *26*, 5015–5030.

Grossberg, S., & Schmajuk, N.A. (1988). Neural dynamics of attentionally modulated Pavlovian conditioning: conditioned reinforcement, inhibition, and opponent processing. In S. Grossberg (Ed.), *Neural networks and natural intelligence*. Cambridge, MA: MIT Press/Bradford Books.

Hallam, B., Hallam, J., & Hayes, G. (1997). A dynamic net for robot control. In O. Omidvar & P. van der Smagt (Eds.), *Neural systems for robotics*. San Diego, CA: Academic Press.

Heikkonen, J., & Koikkalainen, P. (1997). Self-organization and autonomous robots. In O. Omidvar & P. van der Smagt (Eds.), *Neural systems for robotics*. San Diego, CA: Academic Press.

Hilton M., Jawerth B., & Huntsberger T.L. (1993). Local enhancement of compressed images. *Journal of Mathematical Imaging and Vision*, *3*, 39–49.

Huntsberger, T., & Hilton, M. (1995). Distributed multisensor integration in a co-operative multirobot system. In P. S. Schenker & G.T. McKee (Eds.), *Proc. SPIE Symposium on Sensor Fusion and Networked Robotics* (Vol. 2589, pp. 134–141). Philadelphia, PA.

Huntsberger, T.L. (1992a). Data fusion: a neural networks implementation. In M.A. Abidi & R.C. Gonzalez (Eds.), *Data fusion in robotics and machine intelligence*. Orlando, FL: Academic Press.

Huntsberger, T.L. (1992b). Sensor fusion in a dynamic environment. In P.S. Schenker (Ed.), *Proc. SPIE Symposium on Sensor Fusion V* (Vol. 1828, pp. 175–182). Boston, MA.

Huntsberger T.L. (1995). Biologically motivated cross-modality sensory fusion system for automatic target recognition. *Neural Networks*, *8 (7) 8*, 1215–1226.

Huntsberger, T.L. (1996). Hybrid teleoperated/cooperative multirobot system for complex retrieval operations. In P.S. Schenker & G.T. McKee (Eds), *Proc. SPIE Symposium on Sensor Fusion and Distributed Robotic Agents* (pp. 11–17). Boston, MA.

Huntsberger, T.L. (1997). Autonomous multirover system for complex planetary retrieval operations. In P.S. Schenker & G.T. McKee (Eds.), *Proc. SPIE Symposium on Sensor Fusion and Decentralized Control in Autonomous Robotic Systems* (pp. 221–227). Pittsburgh, PA.

Huntsberger, T.L., and Ajjimarangsee, P. (1990). Parallel self-organizing feature maps for unsupervised pattern recognition. *International Journal of General Systems*, *16*, 357–372. Also reprinted in J.C. Bezdek & S.K. Pal (Eds.), *Fuzzy models for pattern recognition* (pp. 483–495). Piscataway, NJ: IEEE Press.

Huntsberger, T.L., & Jawerth, B. (1993). Wavelet based sensor fusion. In P.S. Schenker (Ed.), *Proc. SPIE Symposium on Sensor Fusion VI* (Vol. 2059, pp. 488–498). Boston, MA.

Huntsberger, T.L., & Jawerth, B. (1995). Morphological image analysis within the wavelet representation. In *Proc. SPIE Symposium on Aerospace/Defense Sensing and Dual-Use Photonics* (pp. 1066–1072). Orlando, FL.

Huntsberger, T.L., Jawerth, B., & Chang, J. (1994). Wavelet analysis of fused EPI sequences. In P.S. Schenker (Ed.), *Proc. SPIE Symposium on Sensor Fusion VII* (pp. 157–163). Boston, MA.

Huntsberger, T.L., Rose, J., & Ramaka, A. (1998). Fuzzy-Face: a hybrid wavelet/fuzzy self-organizing feature map system for face processing. *Journal of Biological Systems*. To appear.

Kontoravdis, D., Likas, A., & Stafylopatis, A. (1992). Collision-free movement of an autonomous vehicle using reinforcement learning. In B. Nuemann (Ed.), *ECAI 92* (pp. 666–670). New York: Wiley.

Kröse, B., & van Dam, J. (1997). Neural vehicles. In O. Omidvar & P. van der Smagt (Eds.), *Neural systems for robotics*. San Diego, CA: Academic Press.

Lee T.S. (1996). Image representation using 2D Gabor wavelets. *IEEE Trans. PAMI*, *18 (10)*, 959–971.

Lorenz, K. (1981). *Foundations of ethology*. Berlin: Springer-Verlag.

Maguire E.A., Frackowiak R.S.J., & Frith C.D. (1997). Recalling routes around London: activation of the right hippocampus in taxi drivers. *J. of Neuroscience*, *17 (18)*, 7103–7110.

Mallat S., & Zhong S. (1992). Characterization of signals from multiscale edges. *IEEE Trans. PAMI*, *14*, 710–732.

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

Matarić, M.J. (1992b). Distributed approaches to behavior control. In P.S. Schenker (Ed.), *Proc. SPIE Sensor Fusion V* (Vol. 1828, pp. 373–382). Boston, MA.

Matarić, M.J. (1997). Behavior-based control: examples from navigation, learning, and group behavior. Technical Report. Waltham, MA: Brandeis University.

Matthies L. (1992). Stereo vision for planetary rovers: stochastic modeling to near real-time implementation. *Intern. J. Computer Vision*, *8 (1)*, 71–91.

Miller, D.P. (1992). Reducing software mass through behavior control. In *Proc. SPIE Conf. Cooperative Intelligent Robotics in Space III* (Vol. 1829, pp. 472–475). Boston, MA.

Miller, D.P., Desai, R.S., Gat, E., Loch, J., & Ivlev, R. (1992). Reactive navigation through rough terrain: experimental results. In *Proc. Tenth National Conf. on AI* (pp. 823–828). San Jose, CA.

Moravec, H.P., & Elfes, A. (1985). High resolution maps from wide angle sonar. In *IEEE Conference on Robotics and Automation* (pp. 116–121). Washington, DC.

Parker, L.E. (1994). Heterogeneous multi-robot cooperation. Technical Report 1865. Cambridge, MA: MIT.

Pirjanian, P., & Christensen, H.I. (1997). Behavior coordination using multiple-objective decision making. In P.S. Schenker & G.T. McKee (Eds.), *Proc. SPIE Symposium on Sensor Fusion and Decentralized Control in Autonomous Robotic Systems* (pp. 78–89). Pittsburgh, PA.

Pomerleau D.A. (1991). Efficient training of artificial neural networks for autonomous navigation. *Neural Computation*, *3*, 88–97.

Prescott, T., & Mayhew, J. (1992). Adaptive local navigation. In A. Blake & A. Yuille (Eds.), *Active vision*. Cambridge, MA: MIT Press.

Recce H., & Harris K.D. (1996). Memory for places: a navigational model in support of Marr's theory of hippocampal function. *Hippocampus*, *6 (6)*, 735–748.

Rosenblatt, J.K., & Payton, D.W. (1989). A fine-grained alternative to the subsumption architecture for mobile robot control. In *Proc. IEEE/INNS Joint Conf. on Neural Networks* (pp. 317–324). Washington, DC.

Roumeliotis, S.I., & Bekey, G.A. (1997). An extended Kalman filter for frequent local and infrequent global sensor data fusion. In P.S. Schenker & G.T. McKee (Eds.), *Proc. SPIE Symposium on Sensor Fusion and Decentralized Control in Autonomous Robotic Systems* (pp. 11–22). Pittsburgh, PA.

Samsonovich, A., McNaughton, B.L., Gothard, K., & Skaggs, W.E. (1995). Hippocampus as a path integrator. Poster presented at the XXVth Annual Society for Neuroscience Meeting, San Diego, CA.

Simmons, R., Krotkov, E., & Bares, J. (1991). A six-legged rover for planetary exploration. In *Proc. AIAA Computing in Aerospace 8* (pp. 739–747).

Soltys, M., & Huntsberger, T.L. (1993). Inverse kinematic analysis of articulated figures using the fuzzy self-organizing feature map. In C.H. Dagli, L.I. Burke, B.R. Fernandez & J. Ghosh (Eds.), *Intelligent engineering systems through artificial neural networks* (Vol. 3, pp. 541–546). ASME Press.

Sutherland K.T., & Thompson W.B. (1994). Localizing in unstructured environments: dealing with errors. *IEEE Transactions on Robotics and Automation*, *10 (6)*, 740–754.

Sutton R.S. (1988). Learning to predict by the method of temporal differences. *Machine Learning*, *3*, 9–44.

Tani J., & Fukumura N. (1994). Learning goal-directed sensory-based navigation of a mobile robot. *Neural Networks*, *7 (1)*, 553–564.

Touretzky, D.S., & Redish, A.D. (1995). Landmark arrays and the hippocampal cognitive map. In L. Niklasson & M. Boden (Eds.), *Current Trends in Connectionism—Proceedings of the 1995 Swedish Conference on Connectionism*. Hillsdale, NJ: Erlbaum.

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

Whittaker, R., Bapna, D., Maimone, M., & Rollins, E. (1997). Atacama desert trek: a planetary analog field experiment. In *Proc. Intern. Sympos. on Artificial Intelligence and Automation for Space*. Tokyo, Japan.

Zalama E., Gaudiano P., & Coronado J.L. (1995). A real-time, unsupervised neural network for the low-level control of a mobile robot in a nonstationary environment. *Neural Networks*, *8 (1)*, 103–123