International Journal of Astrobiology 3 (3): 189–207 (2004) Printed in the United Kingdom
DOI: 10.1017/S147355040500220X © 2005 Cambridge University Press

The Cyborg Astrobiologist: first field experience

Patrick Charles McGuire[1,2], Jens Ormö[3], Enrique Díaz Martínez[3,4], José Antonio Rodríguez Manfredi[1], Javier Gómez Elvira[1], Helge Ritter[5], Markus Oesker[5] and Jörg Ontrup[5]

[1]*Robotics Laboratory, Centro de Astrobiología (CAB), Carretera de Torrejón a Ajalvir km 4.5, Torrejón de Ardoz, 28850 Madrid, Spain*
[2]*Transdisciplinary Laboratory, Centro de Astrobiología (CAB), Carretera de Torrejón a Ajalvir km 4.5, Torrejón de Ardoz, 28850 Madrid, Spain*
e-mail: mcguire@physik.uni-bielefeld.de
[3]*Planetary Geology Laboratory, Centro de Astrobiología (CAB), Carretera de Torrejón a Ajalvir km 4.5, Torrejón de Ardoz, 28850 Madrid, Spain*
[4]*Dirección de Geología y Geofísica, Instituto Geológico y Minero de España, Calera 1, Tres Cantos, 28760 Madrid, Spain*
[5]*Neuroinformatics Group, Computer Science Department, Technische Fakultät, University of Bielefeld, PO Box 10 01 31, 33501 Bielefeld, Germany*

Abstract: We present results from the first geological field tests of the 'Cyborg Astrobiologist', which is a wearable computer and video camcorder system that we are using to test and train a computer-vision system towards having some of the autonomous decision-making capabilities of a field-geologist and field-astrobiologist. The Cyborg Astrobiologist platform has thus far been used for testing and development of the following algorithms and systems: robotic acquisition of quasi-mosaics of images; real-time image segmentation; and real-time determination of interesting points in the image mosaics. The hardware and software systems function reliably, and the computer-vision algorithms are adequate for the first field tests. In addition to the proof-of-concept aspect of these field tests, the main result of these field tests is the enumeration of those issues that we can improve in the future, including: detection and accounting for shadows caused by three-dimensional jagged edges in the outcrop; reincorporation of more sophisticated texture-analysis algorithms into the system; creation of hardware and software capabilities to control the camera's zoom lens in an intelligent manner; and, finally, development of algorithms for interpretation of complex geological scenery. Nonetheless, despite these technical inadequacies, this Cyborg Astrobiologist system, consisting of a camera-equipped wearable-computer and its computer-vision algorithms, has demonstrated its ability in finding genuinely interesting points in real-time in the geological scenery, and then gathering more information about these interest points in an automated manner.

Received 30 August 2004, accepted 27 October 2004

Key words: computer vision, co-occurrence histograms, field geology, gypsum, image segmentation, interest map, Mars, Miocene, robotics, uncommon map, wearable computers.

Introduction

> '*With great autonomy comes great responsibility*'
> – with apologies to Peter Parker's Uncle Ben (Lee *et al*. 2002).

Outside of the Mars robotics community, it is commonly presumed that the robotic rovers on Mars are controlled in a time-delayed joystick manner, wherein commands are sent to the rovers several if not many times per day, as new information is acquired from the rovers' sensors. However, inside the Mars robotics community, they have learned that this process is rather cumbersome, and they have developed much more elegant methods for robotic control of the rovers on Mars, with highly significant degrees of robotic autonomy. Particularly, the Mars Exploration Rover (MER) team has demonstrated autonomy for the two robotic rovers Spirit and Opportunity to the level that: practically all commands for a given Martian day (1 'sol' = 24.6 h) are delivered to each rover from Earth before the robot awakens from its power-conserving night-time resting mode (Arvidson *et al*. 2003; Bell *et al*. 2003; Maki *et al*. 2003; Squyres *et al*. 2003; Crisp *et al*. 2003). Each rover then follows the commanded sequence of moves for the entire sol, moving to desired locations, articulating its arm with its sensors to desired points in the workspace of the robot, and acquiring data from the cameras and chemical sensors. From an outsider's point of view, these capabilities may not seem to be significantly autonomous, in that all the commands are being sent from Earth, and the

MER rovers are merely executing those commands. However, upon closer inspection, what the MER team has achieved is truly amazing.

- Firstly, the rovers can move to points 50–150 m away in one sol with autonomous obstacle avoidance enabled for the uncertain or dangerous parts of the journey (Nesnas *et al.* 1999; Goldberg *et al.* 2002).
- Secondly, prior to a given sol, based on information received after the previous sol, the MER team has the remarkable capabilities to develop a command sequence of tens or hundreds of robotic commands for the entire sol. As of 4 July 2004, this was taking 4–5 h per sol for the mission team to complete, rather than the 17 h per sol that it took at the beginning of the MER missions (Squyres 2004b).

Such capabilities for semi-autonomous teleoperated robotic 'movement and discovery' are a significant leap beyond the capabilities of the previous Mars lander missions of Viking I and II and of Pathfinder and Sojourner, particularly in the large improvements in the mobility of both the cameras and the instrumentation onboard the robot, and this allows significant discoveries to be made and new insights to be obtained (Arvidson *et al.* 2004; Bell *et al.* 2004; Catling 2004; Chan *et al.* 2004; Christensen *et al.* 2004; Crumpler *et al.* 2004; Golembek *et al.* 2004; Greeley *et al.* 2004; Herkenhoff *et al.* 2004; McSween *et al.* 2004; Moore 2004; Squyres 2004a; Squyres *et al.* 2004; Squyres & Athena Science Team 2004). Nonetheless, we would like to build upon the great success of the MER rovers by developing enhancing technology that could be deployed in future robotic and/or human exploration missions to the Moon, Mars, and Europa.

One future mission deserves special discussion for the technology developments described in this paper: the Mars Science Laboratory, planned for launch in 2009 (MSL'2009). A particular capability desired for this MSL'2009 mission will be to rapidly traverse to up to three geologically-different scientific points-of-interest within the landing ellipse (Heninger *et al.* 2003). These three geologically-different sites will be chosen from Earth by analysis of relevant remote-sensing imagery from a possible co-launched, satellite-based, Synthetic Aperture Radar (SAR) system or high-resolution maps from prior missions (i.e. 30 cm pixel^{-1} spatial resolution from the HiRISE imager on the Mars Reconnaissance Explorer with imagery expected in 2006–2007 (McEwen *et al.* 2003)). Possible desired maximal traversal rates could range from 300–2000 m sol^{-1} in order to reach each of the three points-of-interest in the landing ellipse in a minimum time. These are our rough estimates of the typical maximum traversal rates, based on current capabilities, mission goals, increased rover size and expectations of the state of technology for the possible 2009 mission; our estimate is partially based on such work as described in Goldberg *et al.* (2002), Olson *et al.* (2003), Crawford (2002), and Crawford & Tamppari (2002).

Given these substantial expected traversal rates of the MSL'2009 rover, autonomous obstacle avoidance (Goldberg *et al.* 2002) and autonomous visual odometry and localization (Olson *et al.* 2003) will be essential to achieve these rates as, otherwise, rover damage and slow science-target approach would occur. Given such autonomy in the rapid traversals, it behoves us to enable the autonomous rover with sufficient scientific responsibility. If we did not, the robotic rover exploration system might drive right past an important scientific target-of-opportunity along the way to the human-chosen scientific point-of-interest. Crawford & Tamppari (2002) and their team summarized possible 'autonomous traverse science', in which science pancam and mini-TES (thermal emission spectrometer) image mosaics are autonomously obtained every 20–30 m during a 300 m traversal (in their example). They stated that 'there *may be* onboard analysis of the science data from the pancam and the mini-TES, which compares this data to predefined signatures of carbonates or other targets of interest. If detected, traverse may be halted and information relayed back to Earth'. This onboard analysis of the science data is precisely the technology issue that we have been working towards solving. This paper is the first detailed published report describing our progress towards giving a robot some aspects of autonomous recognition of scientific targets-of-opportunity (see McGuire *et al.* (2004) for a brief description and motivation of our work).

The general aim of study for this project is to develop autonomous computer vision algorithms for natural terrestrial geological scenery in uncontrolled lighting conditions (see our 'Cyborg Astrobiologist' research platform in Fig. 1), with the expectation that such algorithms will mature in the future so that they can be used in the human or robotic astrobiological exploration of Mars and some of the moons of our Solar System. This project is not only aimed at giving more scientific autonomy to robotic explorers on the surfaces of other planetary bodies, but it is also meant to take some of the burden off of the large team of human geologists at mission control here on the Earth by autonomously performing some of the low-level geological image analysis, thus freeing the human geologists to do more high-level analysis of the scenery. Such a system could also assist human geologists in the field here on the Earth, as well as astronauts in their geological field studies (who may be less-practiced in geology and more burdened with other duties and concerns). Our long-term aim of study is for the Cyborg Astrobiologist system to be able to understand the geological history of an outcrop in order to decide which area should be approached for sampling. Our short-term aim of study is to demonstrate that the Cyborg Astrobiologist system can (in the field) find the most unique regions in an image mosaic, and then to autonomously study those interesting regions in more detail.

This paper is organized as follows. We first give additional introductory background information, then we give a discussion of the hardware and software for our 'Cyborg Astrobiologist' system. Summaries and results of the first two geological expeditions of our Cyborg Astrobiologist system are then presented, followed by a more general discussion and conclusions.

https://doi.org/10.1017/S147355040500220X Published online by Cambridge University Press

(a) (b)

Fig. 1. (a) Díaz Martínez and McGuire with the Cyborg Astrobiologist system at 1:30 p.m. on 3 March 2004, 10 m from the outcrop cliff that was studied during the first geological field mission near Rivas Vaciamadrid. We are taking notes prior to acquiring one of our last-of-the-day mosaics and its set of interest-point chips. (b) A close-up view of the same scene, but focusing on the half-seated/half-tripod Cyborg Astrobiologist and the assisting human geologist. Photo copyright Díaz Martínez, Ormö and McGuire.

Autonomous recognition of scientific astrobiological targets-of-opportunity

We discuss here two of the related efforts in the development of autonomous recognition of scientific targets-of-opportunity for astrobiological exploration: firstly, the work on developing a Nomad robot to search for meteorites in Antarctica led by the Carnegie Mellon University Robotics Institute (CMU-RI) (Apostolopoulos *et al*. 2000; Vandapel *et al*. 2000; Wagner 2000; Pederson 2001); and, secondly, the work by a group at NASA Ames Research Center (NASA-ARC) on developing a Geological Field Assistant (GFA) (Gulick *et al*. 2001, 2002a, b, 2003, 2004). There are several other efforts in the community which also deserve mention including Volpe (2004), Huntsberger *et al*. (2002), Storrie-Lombardi *et al*. (2002), Corsetti & Storrie-Lombardi (2003), Cabrol *et al*. (2001), Wettergreen *et al*. (1999), Whittaker *et al*. (1997), Cheeseman *et al*. (1988), and Cheeseman & Stutz (1996). Outside of the astrobiology community, there is some similar work on automated obstacle avoidance (instead of target discovery) by Batavia & Singh (2001), among others.

The CMU-RI team has programmed the capabilities to autonomously distinguish between rocks and ice/snow in various lighting conditions (direct sunlight, shadows, and diffuse light in overcast conditions) into the science computer of its meteorite-searching Nomad robot (Apostolopoulos *et al*. 2000). This autonomous decision between whether a pixel is ice/snow or if it is rock is decided by how 'blue' the pixel is: if the pixel has a 'blue' ratio (blue/(blue + green)) above a certain lighting-dependent threshold, then it is considered to be ice/snow; otherwise, it is considered to be rock/meteorite. This system works rather well except when the rocks are in shadows (but we naively suggest that a nonlinear separator in the two-dimensional space of 'blue' and intensity could solve this problem). After a rock/meteorite is segmented from the ice/snow based on how 'blue' it is, it can then trigger the Nomad robot's Bayesian classifier network system to compute the probability of the rock being any one of six different meteor types or of 19 different terrestrial rock types, based on data from a metal detector, a colour camera detector, and a spectrometer. The Nomad robot '[autonomously] found and [autonomously] classified five indigenous meteorites [of 42 rocky targets studied and classified over a 2500 m^2 search area] during an expedition to a remote site of Elephant Moraine [in Antarctica] in January 2000' (Apostolopoulos *et al*. 2000).

The NASA-ARC GFA team have developed a number of algorithms for assisting an astronaut in geology activities or in adding capabilities to an autonomous robotic geologist (Gulick *et al*. 2001, 2002a, b, 2003, 2004). They have been building an image and spectral database of rocks and minerals (there are now over 700 such rock and mineral entries in their database) for continued development of advanced classifier algorithms, including future development work to autonomously identify 'key igneous, sedimentary and some metamorphic rocks' (Gulick *et al*. 2003). As of 2003, they were able to identify 'some igneous rocks from texture and colour information; quartz, silica polymorphs, calcite, pyroxene and jarosite from both visible/near-infrared and mid-infrared spectra; and high-iron pyroxenes and iron-bearing minerals using visible/near-infrared spectra only' (Gulick *et al*. 2003). As of 2004, the NASA Jet Propulsion Laboratory (JPL) CLARAty system team (Volpe 2004) had incorporated the NASA-ARC GFA team's rock, layer and horizon detectors (Gulick *et al*. 2001) for 'possible use by MSL'2009 and [other] future Mars robotics missions' (Gulick *et al*. 2004). Work reported in 2004 includes significant progress on their igneous rock detection system: the Decision Tree method has yielded better and more decipherable results than the Bayesian Network approach, giving 'correct identification of greater than 80% of granites and granodiorites and greater than 80% of andesites and basalts, using colour and texture algorithms combined [without spectral information]' (Gulick *et al*. 2004).

The Cyborg geologist and astrobiologist system

Our ongoing effort in the area of autonomous recognition of scientific targets-of-opportunity for field geology and field

https://doi.org/10.1017/S147355040500220X Published online by Cambridge University Press

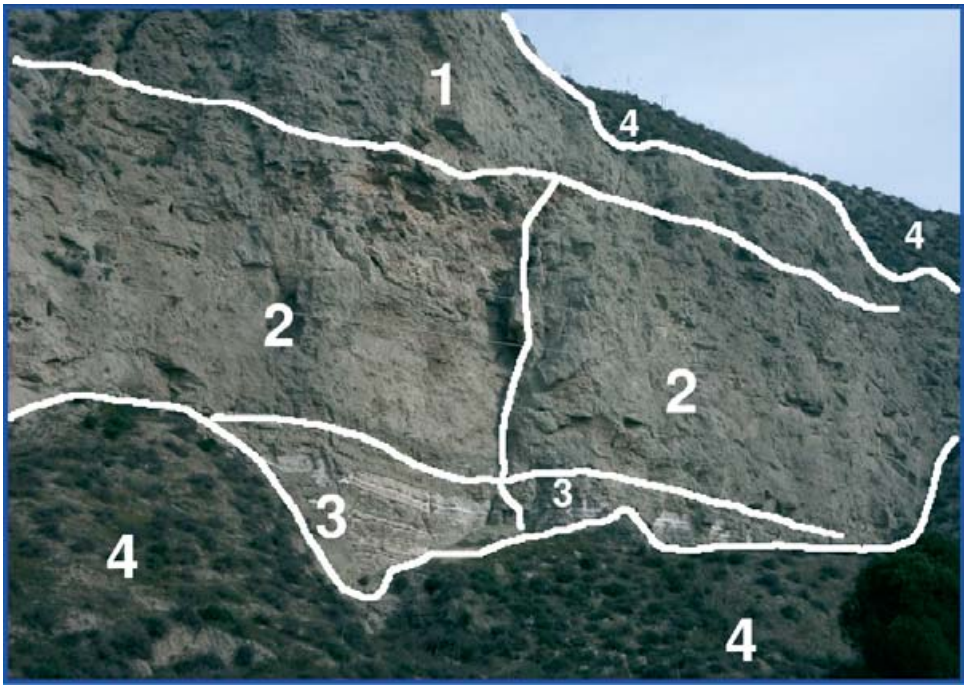

Fig. 2. An image segmentation made by geologist Díaz Martínez of the outcrop during the first mission to Rivas Vaciamadrid. Region 1 has a tan colour and a blocky texture; region 2 is subdivided by a vertical fault and has more red colour and a more layered texture than region 1; region 3 is dominated by white and tan layering; region 4 is covered by vegetation. The dark and wet spots in region 3 were only observed during the second mission, 3 months later. The Cyborg geologist/astrobiologist made its own image segmentations for portions of the cliff face that included the area of the white layering at the bottom of the cliff (see Figs 6–8). Photo copyright Díaz Martínez, Ormö and McGuire.

astrobiology is also beginning to mature. To date, we have developed and field tested a GFA-like 'Cyborg Astrobiologist' system (McGuire *et al.* 2004) that can:

- use human mobility to maneuver to and within geological sites and follow suggestions from the computer as to how to approach a geological outcrop;
- use a portable robotic camera system to obtain a mosaic of colour images;
- use a 'wearable' computer to search in real-time for the most uncommon regions of these mosaic images;
- use the robotic camera system to re-point at several of the most uncommon areas of the mosaic images, in order to obtain much more detailed information about these 'interesting' uncommon areas;
- use human intelligence to choose between the wearable computer's different options for interesting areas in the panorama for closer approach; and
- repeat the process as often as desired, sometimes retracing a step of the geological approach.

In the Mars Exploration Workshop in Madrid in November 2003, we demonstrated some of the early capabilities of our 'Cyborg' geologist/astrobiologist system (McGuire *et al.* 2004). We have been using this Cyborg system as a platform to develop computer-vision algorithms for recognizing interesting geological and astrobiological features, and for testing these algorithms in the field here on the Earth.

The half-human/half-machine 'Cyborg' approach (see Figs 1 and 5) uses human locomotion and human-geologist intuition/intelligence for taking the computer-vision algorithms to the field for teaching and testing, using a wearable computer. This is advantageous because we can therefore concentrate on developing the 'scientific' aspects for autonomous discovery of features in computer imagery, as opposed to the more 'engineering' aspects of using computer vision to guide the locomotion of a robot through treacherous terrain. This means the development of the scientific vision system for the robot is effectively decoupled from the development of the locomotion system for the robot.

After the maturation of the computer-vision algorithms, we hope to transplant these algorithms from the Cyborg computer to the on-board computer of a semi-autonomous robot that will be bound for Mars or one of the interesting moons in our Solar System.

Geological approach to an outcrop: metonymical contacts

This method of using a semi-autonomous system to guide a geological approach to an outcrop, in which the wearable computer decides or gives options as to what in the panorama is the most interesting and deserves closer study, is meant to partially emulate the decision process in the mind of a professional field geologist. An approach from a distance, using information available at further distances to guide the approach to closer distances, is a reasonable, and possibly, generic method used by most field geologists when approaching an outcrop subject to study. Our technique of

https://doi.org/10.1017/S147355040500220X Published online by Cambridge University Press

Fig. 3. For the simple, idealized image on the left, we show the corresponding uncommon map on the right. The whiter areas in the uncommon map are more uncommon than the darker areas in this map.

presenting several point-like options for more close-up study is a reasonable method to 'linearize' the convoluted thought processes of a practicing field geologist.

McGreevy (1992, 1994) has completed significant detailed cognitive and ethnographic studies of the practice of field geology, using a head-mounted camera and virtual-reality visor system for part of his study as well as extensive on-site interviews of field-geologist colleagues. One particularly interesting concept from McGreevy's work is that of 'metonymic' geological contacts. For field-geology, this concept of metonymy is the juxtaposition of differing geological units as well as the relation between them; metonymy is to differences, as metaphor is to similarities.

Both of the field geologists on our team have also independently stressed the importance to field geologists of such geological contacts and the differences between the geological units that are separated by the geological contact. For this reason, in March 2003, the roboticist and the field geologists decided that the most important tool at the beginning of our computer-vision algorithm development was that of 'image segmentation'. Such image segmentation algorithms would allow the computer to break down a panoramic image into different regions (see Fig. 2 for an example), based on similarity, and to find the boundaries or contacts between the different regions in the image, based on difference. Much of the remainder of this paper discusses the first geological field trials with the wearable computer of the segmentation algorithm we have developed over the last year.

Image segmentation, uncommon maps, interest maps, and interest points

With human vision, a geologist:

- first tends to pay attention to those areas of a scene which are most unlike the other areas; and
- then attempts to find the relation between the different areas of the scene, in order to understand the geological history of the outcrop.

The first step in this prototypical thought process of a geologist was our motivation for inventing the concept of uncommon maps. See Fig. 3 for a simple illustration of the concept of an uncommon map. We have not yet attempted to solve the second step in this prototypical thought process of a geologist, but it is evident from the formulation of the second step, that human geologists do not immediately ignore the common areas of the scene. Instead, human geologists catalogue the common areas and put them in the back of their minds for 'higher-level analysis of the scene', or in other words, for determining explanations for the relations of the uncommon areas of the scene with the common areas.

Prior to implementing the 'uncommon map', the first step of the prototypical geologist's thought process, we needed a segmentation algorithm in order to produce pixel-class maps to serve as input to the uncommon map algorithm. We have implemented the classic co-occurrence histogram algorithm (Haralick *et al.* 1973; Haddon & Boyce 1990). For this work, we have not included texture information in the segmentation algorithm and the uncommon map algorithm. Currently, each of the three bands of H,S,I colour information is segmented separately (see the software section below and the appendix), and later merged in the interest map by summing three independent uncommon maps. In future work, image segmentation using colour and texture simultaneously could be developed and tested on the Cyborg Astrobiologist system (e.g. Freixenet *et al.* 2004).

The concept of an 'uncommon map' is our invention, although it has probably been independently invented by other authors, since it is somewhat useful. In our implementation, the uncommon map algorithm takes the top eight pixel classes determined by the image segmentation algorithm, and ranks each pixel class according to how many pixels there are in each class. The pixels in the pixel class with the greatest number of pixel members are numerically labelled as 'common', and the pixels in the pixel class with the least number of pixel members are numerically labelled as 'uncommon'. The 'uncommonness' hence ranges from 1 for a common pixel to 8 for an uncommon pixel, and we can therefore construct an uncommon map given any image segmentation map. Rare pixels that belong to a pixel class of 9 or greater are usually noise pixels in our tests thus far, and so are currently ignored. In our work, we construct several uncommon maps from the colour image mosaic, and then we sum these uncommon maps together in order to arrive at a final interest map. This summing to determine an interest map has been also studied by Rae *et al.* (1999). In our implementation for this paper, the interest map is a convenient name to call the sum of several uncommon maps,

https://doi.org/10.1017/S147355040500220X Published online by Cambridge University Press

so the interest map is also an uncommon map in a sense. However, in general, the interest map could also sum other maps and in the implementation of Rae *et al.* (1999), their 'cortical' interest map algorithm not only sums several different types of maps, but the algorithm also adapts the weighting coefficients for these maps in order to highlight the summand map(s) which are most salient, most uncommon, or most interesting. For our work, we have not activated this adaptive summing of the interest map algorithm and we only use uncommon maps in the interest map sum.

In this paper, we develop and test a simple, high-level concept of interest points of an image, which is based on finding the centroids of the smallest (most uncommon) regions of the image. Such a 'global' high-level concept of interest points differs from the lower-level 'local' concept of interest points (Förstner 1986; Förstner & Gülch 1987) based on corners and centers of circular features. However, this latter technique with local interest points is used by the MER team for their stereo-vision image matching and for their visual-odometry and visual-localization image matching (Nesnas *et al.* 1999; Goldberg *et al.* 2002; Arvidson *et al.* 2003; Olson *et al.* 2003). Our interest point method bears somewhat more relation to the higher-level wavelet-based salient points technique (Sebe *et al.* 2003), in that they search first at coarse resolution for the image regions with the largest gradient, and then use wavelets in order to zoom in towards the salient point within that region that has the highest gradient. Their salient point technique is edge-based, whereas our interest point is currently region-based. Since, in the long term, we have an interest in geological contacts, this edge-based and wavelet-based salient point technique could be a reasonable future interest-point algorithm to incorporate into our Cyborg Astrobiologist system for testing.

For a review of some of the general aspects of computer vision for natural scenery see Batlle *et al.* (2000). We give a description of our specific software system in the software section below and more details are given in the appendix.

Hardware for the Cyborg Astrobiologist

The non-human hardware of the Cyborg Astrobiologist system consists of:

- a 667 MHz wearable computer (from ViA Computer Systems in Minnesota) with a 'power-saving' Transmeta 'Crusoe' CPU and 112 MB of physical memory;
- a SV-6 Head Mounted Display (from Tekgear in Virginia, via the Spanish supplier Decom in València) with native pixel dimensions of 640×480 that works well in bright sunlight;
- a SONY 'Handycam' colour video camera (model DCR-TRV620E-PAL);
- a thumb-operated USB finger trackball from 3G Green Green Globe Co., resupplied by ViA Computer Systems and by Decom;
- a small keyboard attached to the human's arm;
- a tripod for the camera; and
- a pan-tilt unit (model PTU-46-70W) from Directed Perception in California with a bag of associated power and signal converters.

The power-saving aspect of the wearable computer's Crusoe processor is important because it extends battery life, meaning that the human does not need to carry very many spare batteries, meaning less fatigue for the human geologist. A single lithium-ion battery, which weighs about 1 kg, can last for 3 h or so for this application. The SONY Handycam provides real-time imagery to the wearable computer via an IEEE1394/Firewire communication cable. The current system uses five independent batteries, one for each sub-system: the computer, the head-mounted display, the camera, the pan-tilt unit, and the communication interface between the pan-tilt unit and the computer. Assuming fully-charged batteries, all of the batteries are currently more than sufficient for a mission of several hours.

The main reason for using the tripod is that it allows us to use the pan-tilt unit for mosaicking and for re-pointing. Without the tripod, there would be an unusual amount of jitter in the image-registration caused by the human fidgeting with the camera, regardless of whether the camera is hand-mounted, shoulder-mounted or head-mounted. In future studies, if the mosaicking or re-pointing is not desired, then neither the pan-tilt unit nor the tripod would be necessary; eliminating both the tripod and the pan-tilt unit would certainly add to the mobility of the Cyborg Astrobiologist system.

For this particular study, which included the pan-tilt unit and the tripod, the wearable computer could have been replaced by a suitable non-wearable computer. However, we would like to retain flexibility in our work for future studies, in which the mobility granted by the wearable computer may be essential. The head-mounted display was much brighter than our earlier tablet display, and together with the thumb-operated finger mouse, the head-mounted display was more ergonomic than our earlier tablet display and stylus. However, the spatial resolution of the head-mounted display was somewhat less than the resolution of the tablet display, but sufficient for the current task. It would be useful to be able to switch rapidly between the tablet display and the head-mounted display, but our system is not capable of this at this moment.

The thumb-operated finger trackball worked rather well for navigating the Microsoft Windows environment of the wearable system. However, when used outdoors in bright sunlight, we often had to shade the trackball with the user's second hand in order for it to function properly. Since this mouse is an optical mouse, the mouse fails when sufficient sunlight leaks into the trackball. Alternatively, we could have sealed the seams in the plastic mouse with tape or glue. Additionally, since the wearable computer only has one USB port, the system would have been easier to use when downloading data to a USB memory stick if the USB finger trackball had been the PS2 model of the finger trackball; this will be changed before the next mission.

https://doi.org/10.1017/S147355040500220X Published online by Cambridge University Press

Two other issues affected the usage and mobility of the wearable system in the field. First, the serial port of the ViA wearable computer often failed to provide a signal of sufficient amplitude; this signal is needed to drive the pan-tilt unit used for robotic exploration of the environment. We sent the wearable computer system back to the manufacturer for investigation, and they suggested that we reinstall the COM port in software when we had this problem. However, since this problem was so frequent (especially after putting the wearable computer into sleep mode for power conservation), we decided to fix the signal problem with an external operational amplifier (OpAmp). This external OpAmp system functions well, but it requires occasional 9 V battery changes, and it needs a certain amount of engineering improvements in ergonomics and miniaturization. Second, the pan-tilt unit requires +12 to +30 V of power. The battery we have chosen for this task weighs about 2 kg and lasts for 5–10 h with the currently-used Cyborg software. The battery, together with the signal conversion box (from serial signals to PTU signals) and the signal OpAmp box, require an extra, small carrying bag in order to use the PTU in the field. This bag is usually placed on the ground when each tripod position is reached, but for transport between the tripod positions the PTU bag is added baggage and puts some minor limits on mobility. The PTU system certainly merits further engineering.

Software for the Cyborg Astrobiologist

The wearable computer processes the images acquired by the colour digital video camera, to compute a map of interesting areas. The computations include simple mosaicking by image-butting, as well as two-dimensional histogramming for image segmentation (Haralick *et al.* 1973; Haddon & Boyce 1990). This image segmentation is independently computed for each of the Hue, Saturation, and Intensity (for *H*,*S*,*I*) image planes, resulting in three different image-segmentation maps. These image-segmentation maps were used to compute 'uncommon' maps (one for each of the three *H*,*S*,*I* image-segmentation maps): each of the three resulting uncommon maps gives highest weight to those regions of smallest area for the respective *H*,*S*,*I* image planes. Finally, the three *H*,*S*,*I* uncommon maps are added together into an interest map, which is used by the Cyborg system for subsequent interest-guided pointing of the camera. The image-processing algorithms and robotic systems-control algorithms are all programmed using the graphical programming language NEO/NST (Ritter *et al.* 1992, 2002). Using such a graphical programming language adds certain flexibility and ease-of-understanding to our Cyborg Astrobiologist project, which is, by its nature, largely a software project. We discuss some of the details of the software implementation in the appendix.

After segmenting the mosaic image (see the appendix and Fig. 9), it becomes obvious that a very simple method to find interesting regions in an image is to look for those regions in the image that have a significant number of uncommon pixels. We accomplish this by (see Fig. 10) for an example) first creating an uncommon map based on a linear reversal of the segment peak ranking (we have also studied nonlinear reversals of the segment ranking). We then add the three uncommon maps (for *H*,*S*,*I*) together to form an interest map. We finish by blurring this interest map.[1]

Our study extends Rae, Fislage and Ritter's interest map concept by using uncommon maps based on image segmentation of *H*,*S*,*I* for the component maps of the interest map, rather than using the *H*,*S*,*I* maps themselves for the component maps of the interest map. This extension is especially important for our purposes of finding the most uncommon and interesting areas of an image. By using such an image-fusion technique as Rae, Fislage and Ritter's cortical interest map technique to combine three uncommon/segmentation maps we also avoided the complexity of developing a true image-segmentation algorithm for *colour* images in order to determine a single uncommon/segmentation map. If the current work is sufficiently valued by the astrobiology or other communities, then maybe the Cyborg platform could be used for the development, testing, and optimization of true colour image-segmentation algorithms.

Based on the three largest peaks in the blurred/smoothed interest map, the Cyborg system then guides the pan-tilt unit to point the camera at each of these three positions to acquire high-resolution colour images of the three interest points, as shown in four different examples in Figs 6–8. The robotic control and image mosaicking portion of the Cyborg software required significant time for development and debugging since the logic of this autonomy is somewhat complicated. However, by extending a simple image-acquisition and image-processing system to include robotic and mosaicking elements, we were able to conclusively demonstrate that the system can make reasonable decisions by itself in the field for robotic pointing of the camera (see Figs 6–8).

Descriptive summaries of the field site and expeditions

On 3 March 2004, three of the authors (McGuire, Díaz Martínez and Ormö) tested the 'Cyborg Astrobiologist' system for the first time at a geological site, the gypsum-bearing southward-facing stratified cliffs near the 'El Campillo' lake of Madrid's Southeast Regional Park (IGME 1975; Calvo *et al.* 1989; Castilla Cañamero 2001), outside the suburb of Rivas Vaciamadrid.

After post-mission analysis of the results from the Cyborg Astrobiologist's first mission, a second mission was planned for 12 May 2004. After problems were discovered in preparation for the second mission, we decided to delay the second mission to Rivas Vaciamadrid until the 11 June 2004. For the second rescheduled mission, McGuire and Ormö took the Cyborg Astrobiologist system to the same southward-facing stratified cliffs outside of Rivas. Due to the

[1] We blurred with a Gaussian smoothing kernel of width $B = 10$. This smoothing kernel effectively gives more weight to clusters of uncommon pixels, rather than to isolated, rare pixels.

https://doi.org/10.1017/S147355040500220X Published online by Cambridge University Press

significant storms in the previous 3 months, there were two points in the gypsum cliffs that were leaking water down the face of the cliff (despite the hot weather before and during the second mission).

Description and choice of the field site at Rivas Vaciamadrid

The rocks at the outcrop are of Miocene age (15–23.5 Myrs old), and mainly consist of gypsum and clay minerals, with other minor minerals present (anhydrite, calcite, dolomite, etc.). These rocks form the whole cliff face that we studied, and were all deposited in and around lakes with high evaporation (evaporitic lakes) during the Miocene. They belong to the so-called Lower Unit of the Miocene. Above the cliff face is a younger unit (which we did not study) with sepiolite, chert, marls, etc., forming the white relief seen above from the distance, which is part of the so-called Middle Unit or Intermediate Unit of the Miocene. The colour of the gypsum is normally greyish, and in-hand specimens of the gypsum contain large and clearly visible crystals.

We chose this locality for several reasons:

- it is a so-called 'outcrop', meaning that it is exposed bedrock (without vegetation or soil cover);
- it has distinct layering and significant textures, which will be useful for this study and for future studies;
- it has some degree of colour differences across the outcrop;
- it has tectonic offsets, which will be useful for future high-level analyses of searching for textures that are discontinuously offset from other parts of the same texture;
- it is relatively close to our workplace;
- it is not a part of a noisy and possibly dangerous construction site or highway site.

The upper areas of the chosen portion of the outcrop at Rivas Vaciamadrid (see Fig. 2) were mostly of a tan colour and a blocky texture. There was some faulting in the more layered middle areas of the outcrop, accompanied by some slight red colouring. The lower areas of the outcrop were dominated by white and tan layering. Due to some differential erosion between the layers of the rocks, there was some shadowing caused by the relief between the layers. However, this shadowing was perhaps at its minimum for direct lighting conditions, since both expeditions were taken at midday and since the outcrop runs from East to West. Therefore, by performing our field-study at Rivas during midday, we avoided longer shadows that might occur at dawn or dusk.

Díaz Martínez had studied similar geological sites in the region of Rivas Vaciamadrid prior to the missions. He explained the bright white area at the bottom of the cliffs as being due to mineral efflorescence (most probably calcium and magnesium sulphates) on the *surface* of the cliff rock. This cliff rock is mostly gypsum and clays. The efflorescence is just a surface feature formed by evaporation of the water seeping out from the rock pores by capillarity.

Details from the first mission to Rivas

We arrived at the site at 10:30 a.m. on 3 March 2004, and for the next hour the two geologists of our team both made independent expert-geologist assessments of the area. In their

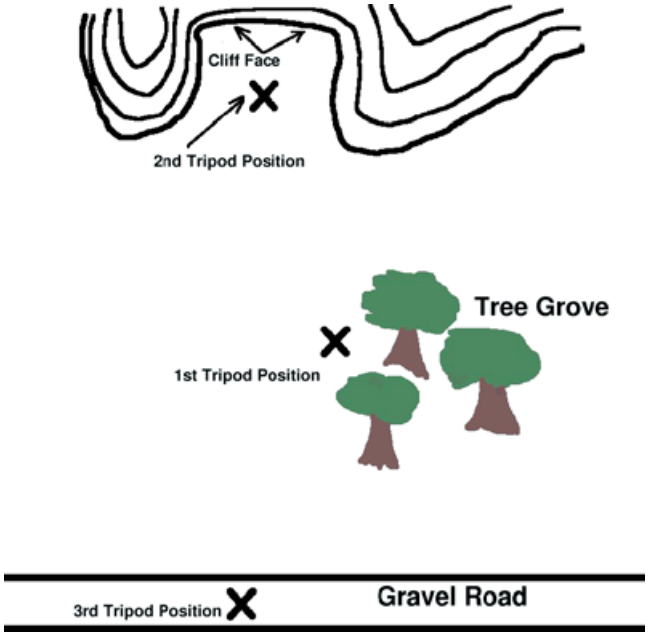

Fig. 4. Map of the three tripod positions of the Cyborg Astrobiologist system during the second mission to Rivas Vaciamadrid (this map is not to scale). See the section on the second mission to Rivas for a description of the geological approach suggested by the Cyborg Astrobiologist to the outcrop. Figure 1 shows an image of the system in action during the first mission, when it was near the second tripod position shown here. Figure 5 shows an image of the system in action during the first mission when it was near the first tripod position shown here. Figures 6–8 show mosaics obtained during the second mission for the third, first and second tripod positions shown here, respectively, and the higher resolution colour imagery of the top three interest points for each of the image mosaics.

assessments, these two geologists included their reasoning processes of how they would gradually approach and study the site, from an initial distance of about 300 m down to a final distance of physical contact with the geological structures. An analysis of the geologists' assessments will be made available in the near future. However, in Fig. 2 we show the segmentation of the outcrop, according to the geologist Díaz Martínez, for reference and as a sample of the geologists' assessments.

At about 11:30 a.m. on the same day, the roboticist finished setting up the Cyborg Astrobiologist system, and started taking mosaic images of the cliff face, with the computer-controlled and motorized colour video camera. It generally takes some time to set up the system (in this case one hour) because:

- several cables need to be attached and secured;
- the computer needs to boot;
- three programs need to be activated and checked;
- communication between the serial port and the pan-tilt unit needs to be confirmed; and
- the non-computer-controlled zoom factor on the camera-lens needs to be set and optimized for the geological scenery, and the pan-tilt unit's step size needs to be set to match the camera's zoom factor.

https://doi.org/10.1017/S147355040500220X Published online by Cambridge University Press

(a) (b)

Fig. 5. (a) Astrobiologist and geologist Ormö is wearing the computer part of the Cyborg Astrobiologist system, as astrobiologist and roboticist McGuire looks on. (b) The team of astrobiologists together for a group pose; during both of these pictures, the robotic part of the Cyborg Astrobiologist was slaving away, acquiring and processing the images of a 9×4 vertically oriented mosaic. These pictures were taken during the 3 March 2004 expedition to Rivas Vaciamadrid. Note the absence of the black spots near the bottom of the cliff face; these black spots appeared some time after this first mission and before the second mission. Photos copyright Díaz Martínez, Ormö and McGuire.

The computer was worn on the user's belt, and typically took 3–5 min to acquire and compose a mosaic image composed of $M \times N$ subimages. Typical values of $M \times N$ used were 3×9 and 11×4. The sub-images were downsampled in both directions by a factor of 4–8 during these tests; the original sub-image dimensions were 360×288.

Several mosaics were acquired of the cliff face from a distance of about 300 m, and the computer automatically determined the three most interesting points in each mosaic. Then, the wearable computer automatically repointed the camera towards each of the three interest points, in order to acquire *non-downsampled* colour images of the region around each interest point in the image. All the original mosaics, derived mosaics and interest-point subimages were then saved to hard disk for post-mission study. During the entire expedition, the robotics engineer could observe what the wearable computer system was doing through a very bright, sunlight-readable head-mounted computer display (HMD). The robotics engineer could and did make adjustments on-the-fly to the automated-exploration settings and algorithms by using this HMD in conjunction with a hand-held free-movement thumbmouse and a small arm-mounted keyboard.

Three other tripod positions were chosen for acquiring mosaics and interest-point image-chip sets. At each of the three tripod positions, two or three mosaic images and interest-point image-chip sets were acquired. One of the chosen tripod locations was about 60 m (see Fig. 5) from the cliff's face and the other two were both about 10 m (see Fig. 1) from the cliff face. There was some concern that if we were to get much closer for any significant amount of time, then rocks from the fragile gypsum-bearing cliffs could fall on us (there was plenty of evidence of past rock falls). We kept a distance, but noticed no such rock falls during the time we were there.

Details from the second mission to Rivas

Ormö and McGuire arrived at 12:00 p.m. on 11 June 2004. One of us first suspected that somebody had lit some fires near the bottom of the cliffs, since there were two black stains, which were medium-sized and soot-like, covering in total about 10–15% (about 4–6 m[2]) of the bright white area at the bottom of the cliffs. Closer analysis discounted this suspicion, as the black stains were really caused by two separate leaks of water[2] from the cliff face, with each leak wetting and darkening a small region of the cliff face.

Due to the heat of the sun, we chose our first Cyborg Astrobiologist tripod position (see the map in Fig. 4) to be in the cool shade of the grove of trees about 60 m from the cliff face (see Fig. 5). We stayed at this tripod position for about 1.3 h, acquiring several 2×2 mosaics, and several 4×3 mosaics, with a downsampling of 4. The system most often determined the wet spots (Fig. 7) to be the most interesting regions on the cliff face. This was encouraging to us, because we also found these wet spots to be the most interesting regions[3].

To end the mission, we chose two other tripod positions (see the map in Fig. 4) for the Cyborg Astrobiologist to acquire mosaics, to search for interest points and then to record detailed information about the points. The first of these two tripod positions was about 10 m from the cliff face, which was significantly closer than the original tripod position in the grove of trees and was approximately the same position as in

[2] Audible from a distance of 60 m.

[3] These dark and wet regions were interesting to us partly because they give information about the development of the outcrop. Even if the relatively small spots were only dark, and not wet (i.e. dark dolerite blocks or a brecciated basalt), their uniqueness in the otherwise white and tan outcrop would have drawn our immediate attention. Additionally, even if this had been our first trip to the site and if the dark spots had been present during this first trip, these dark regions would have captured our attention for the same reasons. The fact that these dark spots had appeared after our first trip and before the second trip was not of paramount importance to grab our interest (but the 'sudden' appearance of the dark spots between the two missions did arouse our higher-order curiosity).

https://doi.org/10.1017/S147355040500220X Published online by Cambridge University Press

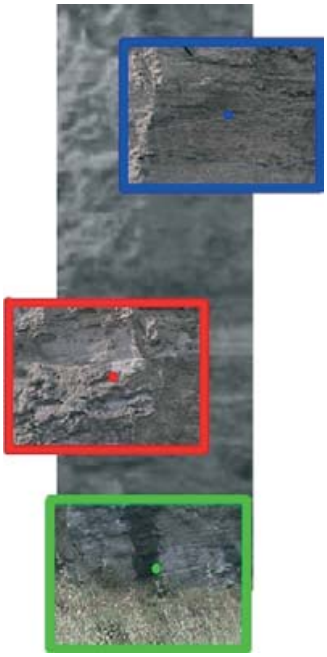

Fig. 6. Mosaic image of a four-by-one set of greyscale sub-images acquired by the Cyborg Astrobiologist at the end of its second expedition (although we should have acquired it at the beginning of the mission). The three most interesting points were subsequently revisited by the camera in order to acquire full-colour higher-resolution images of these points-of-interest. The coloured points and rectangles represent the points that the Cyborg Astrobiologist determined (on location) to be most interesting: *green* is most interesting, *blue* is second most interesting, and *red* is third most interesting. The images were taken and processed in real-time between 4:59 p.m. and 5:02 p.m. local time on 11 June 2004 about 300 m from some gypsum-bearing southward-facing cliffs near the 'El Campillo' lake of the Madrid southeast regional park outside of Rivas Vaciamadrid. The most interesting point (*green* box) corresponds to the wet and dark area near the bottom of the cliffs; this interest point could have been used to guide the geological approach to 30 m and then to 10 m to the interesting dark and wet spot by the Cyborg Astrobiologist earlier in the day, as represented in the next two figures.

the first mission, as depicted in Fig. 1. During this 'close-up' study of the cliff face, we intended to focus the Cyborg Astrobiologist exploration system on the two points that it found most interesting when it was in the more distant tree grove, namely the two wet and dark regions of the lower part of the cliff face. By moving from a distance of 60 m to a distance of 10 m and by focusing at the closer distance on the interest points determined at the larger distance, we wished to simulate how a truly autonomous robotic system would approach the cliff face. Unfortunately, due to a combination of a lack of human foresight in the choice of tripod position and a lack of more advanced software algorithms to mask out the surrounding and less-interesting region (see the discussion in the results section), for one of the two dark spots (Fig. 8(a)), the Cyborg system only found interesting points on the undarkened periphery of the dark and wet stains. Furthermore, the other dark spot was spatially complex, being subdivided into several regions, with some green and brown foliage covering part of the mosaic (Fig. 8(b)). Therefore, in both close-up cases the value of the interest mapping is debatable. This interest mapping could be improved in the future, as we discuss in the section on the results from the second geological field test.

The third and final tripod position was at a distance of about 300 m, near our vehicle on the road leading to the cliffs. We were careful to choose a mosaic that properly masked out the grass and trees as well as the sky. This mosaic ended up being a 1×4 mosaic (Fig. 6) with a downsampling of 8, which spanned the 10 m width of the hollow in the cliffs up to a height of about 30 m. Again, we were pleased when the system found the dark and wet spot near the bottom of the cliff face the most interesting point. If we had begun the Cyborg Astrobiologist's second mission from this third tripod position, instead of immediately going to the shady tree grove, then the sum of these results (from the third tripod position to the first tripod position and either discounting or improving the second tripod position) would demonstrate that the computer-vision algorithms inside the Cyborg Astrobiologist could have served as very reasonable guidance for the geological approach of an autonomous robot towards the cliff face.

Results

Results from the first geological field test

As first observed during the first mission to Rivas on 3 March 2004, the characteristics of the southward-facing cliffs at Rivas Vaciamadrid consist of mostly tan-coloured surfaces, with some white veins or layers, and with significant shadow-causing three-dimensional structures. The computer-vision algorithms performed adequately for a first visit to a geological site, but they need to be improved in the future. At the end of the first mission, the mission team decided that the improvements should include:

- shadow-detection and shadow-interpretation algorithms;
- segmentation of the images based on microtexture; and
- better methods to control the adaptation of the weighting coefficients determined by the cortical interest map algorithm (Rae *et al.* 1999), which we used to determine our interest map.

In the last case, we decided that due to the very monochromatic and slightly shadowy nature of the imagery, the cortical interest map algorithm non-intuitively decided to concentrate its interest on differences in intensity, and tended to ignore hue and saturation.

After the first geological field test, we spent several months studying the imagery obtained during this mission, and fixing various further problems that were discovered during and after the first mission. Although we had hoped that the first mission to Rivas would have been more like a science mission, in reality it was more of an engineering mission.

https://doi.org/10.1017/S147355040500220X Published online by Cambridge University Press

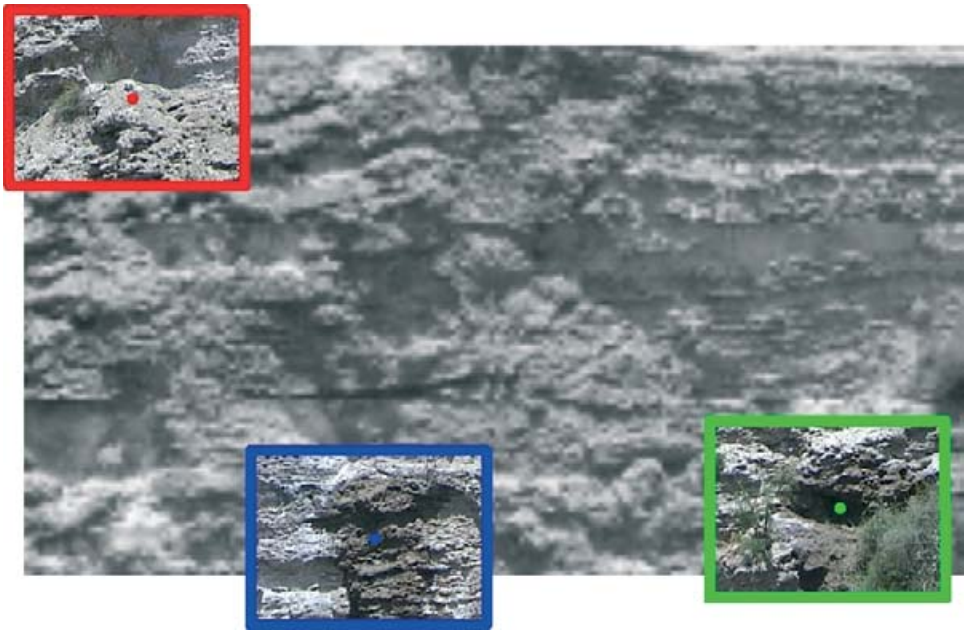

Fig. 7. Mosaic image of a 3×4 set of greyscale sub-images acquired by the Cyborg Astrobiologist at the beginning of its second expedition. The three most interesting points were subsequently revisited by the camera in order to acquire full-colour higher-resolution images of these points-of-interest. The coloured points and rectangles represent the points that the Cyborg Astrobiologist determined (on location) to be most interesting: *green* is most interesting, *blue* is second most interesting, and *red* is third most interesting. The images were taken and processed in real-time between 1:25 p.m. and 1:35 p.m. local time on 11 June 2004 about 60 m from some gypsum-bearing southward-facing cliffs near the 'El Campillo' lake of the Madrid southeast regional park outside of Rivas Vaciamadrid. See Figs 9 and 10 for some details about the real-time image processing that was done in order to determine the location of the interest points in this figure.

Results from the second geological field test

In Figs 6–8, we show the original mosaics of the intensity sub-images for the three tripod positions from the second mission to Rivas Vaciamadrid. The computer then found the three most interesting points and repointed the camera at each of these three points, in order to acquire higher-resolution colour images. This is an emulation of a mapping and exploration camera being used first to get the 'lay of the land', after which a higher technology science instrument (i.e. a mapping thermal emission spectrometer or Mössbauer spectrometer) is trained for more time or with more resources on the most interesting points in the broad landscape. Figures 9 and 10 show how the computer arrived at its final decision as to where the most interesting points are located, for one of these mosaic images.

In Fig. 6, at a distance of 300 m, the system found the dark wet spot at the bottom of the 4×1 mosaic to be the most interesting, the rough texture at the top of the mosaic to be the second most interesting, and a patchy white and tan coloured region in the middle of the mosaic to be the third most interesting. The uncommon map for this mosaic (not shown) shows that the system found the dark and wet spot to be most interesting, largely due the preponderance of pixels of uncommon intensity in this spot. There was also the significant influence of uncommon saturation pixels for this spot. Furthermore, the uncommonness of the saturation pixels in the dark wet spot (and also the second rough texture interest point) is due to the fact that these saturation pixels are outside the common central peak of the saturation histogram. For our pixel-segmentation algorithm, when the pixels are outside the common central peak, these pixels are often assigned to be in one of the uncommon segments. Figure 6 might suggest that the algorithm found that the telephone cable (in front of the rough texture at the top of the mosaic) to be the second most interesting point. However, we find that the rough saturation texture is itself responsible for this region being interesting, and not the telephone cable as one might suspect.

In Fig. 7, from the tree grove at a distance of 60 m, the Cyborg Astrobiologist system found the dark and wet spot on the right-hand side to be the most interesting, the dark and wet spot on the left-hand side to be the second most interesting, and the small dark shadow in the upper left-hand corner to be the third most interesting. For the first two interest points (the dark and wet spots), it is apparent from the uncommon map for intensity pixels in Fig. 10 that these points are interesting due to their relatively remarkable intensity values. By inspection of Fig. 9, we see that these pixels which reside in the white segment of the intensity segmentation mosaic are unusual because they are a cluster of very dim pixels (relative to the brighter segments, with red, blue and green segment-label colours). Within the dark wet spots, we observe that these particular points in the white segment of the intensity segmentation in Fig. 9 are interesting because they reside in the *shadowy* areas of the dark and wet spots. We interpret the interest in the third interest point to be due to the juxtaposition of the small green plant with the shadowing in

https://doi.org/10.1017/S147355040500220X Published online by Cambridge University Press

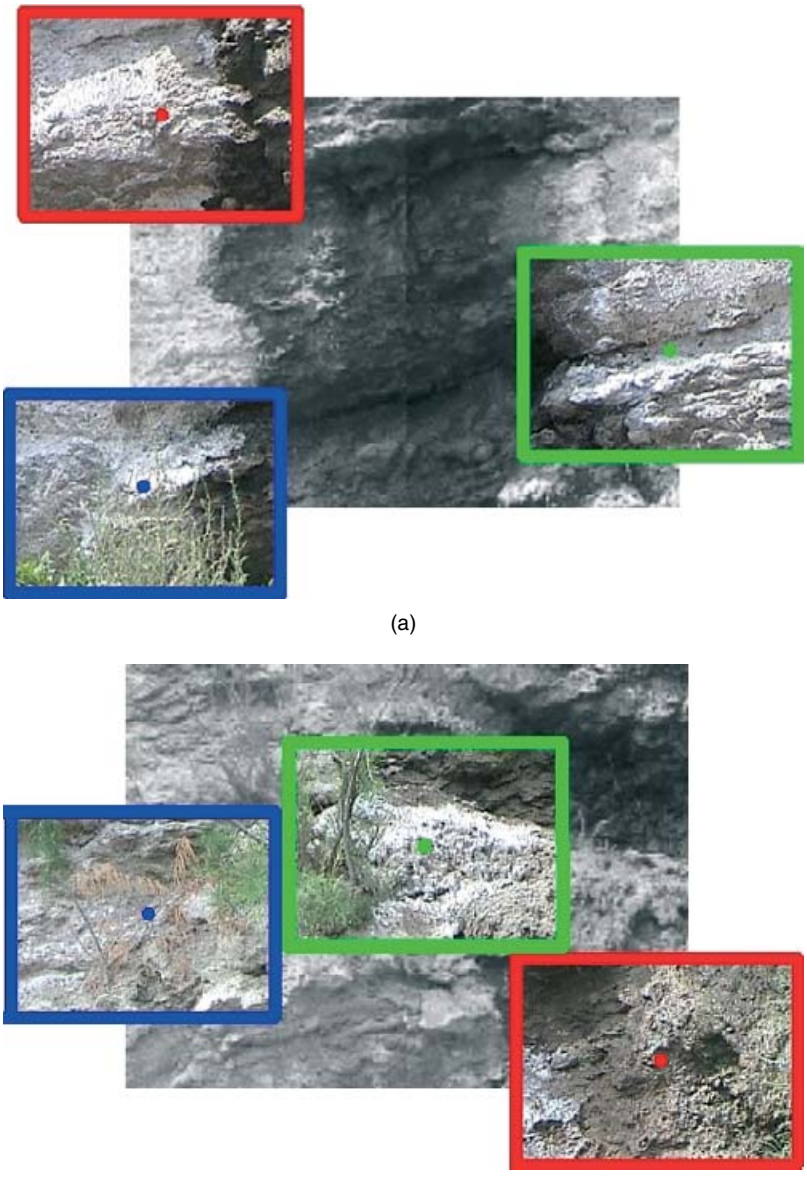

Fig. 8. Two different mosaic images of a two-by-two set of greyscale sub-images acquired by the Cyborg Astrobiologist in the middle of its second expedition. The three most interesting points were subsequently revisited by the camera in order to acquire full-colour higher-resolution images of these points-of-interest. The coloured points and rectangles represent the points that the Cyborg Astrobiologist determined (on location) to be most interesting: *green* is most interesting, *blue* is second most interesting, and *red* is third most interesting. The images were taken and processed in real-time between 3:27 p.m. and 4:05 p.m. local time on 11 June 2004 about 10 m from some gypsum-bearing southward-facing cliffs near the 'El Campillo' lake of the Madrid southeast regional park outside of Rivas Vaciamadrid. The upper mosaic (a) represents the Cyborg's study of the left-most dark and wet spot of Fig. 7, and the lower mosaic (b) represents the Cyborg's study of the right-most dark and wet spot of Fig. 7.

this region; the interest in this point is significantly smaller than for the other two interest points.

A similar analysis of the segmentation, uncommon and interest maps (not shown here) corresponding to Figs 8(a) and 8(b), shows (for example):

- the rareness of the two most interesting green- and blue-boxed regions on the periphery of the dark and wet spot in Fig. 8(a) is due to the fact that they are particularly bright compared with most of the other pixels;
- the peculiarity of the most interesting central white region in Fig. 8(b) is due to the fact that it is particularly bright compared with most of the other pixels.

For both of these cases, the interesting bright regions seem to have a frothy white crust, more so than the nearby regions which are more tan coloured.

The wearable computer's analysis of the close-up 2×2 mosaics of the dark and wet spots (shown in Figs 8(a) and 8(b)) was somewhat disappointing to us, particularly its

https://doi.org/10.1017/S147355040500220X Published online by Cambridge University Press

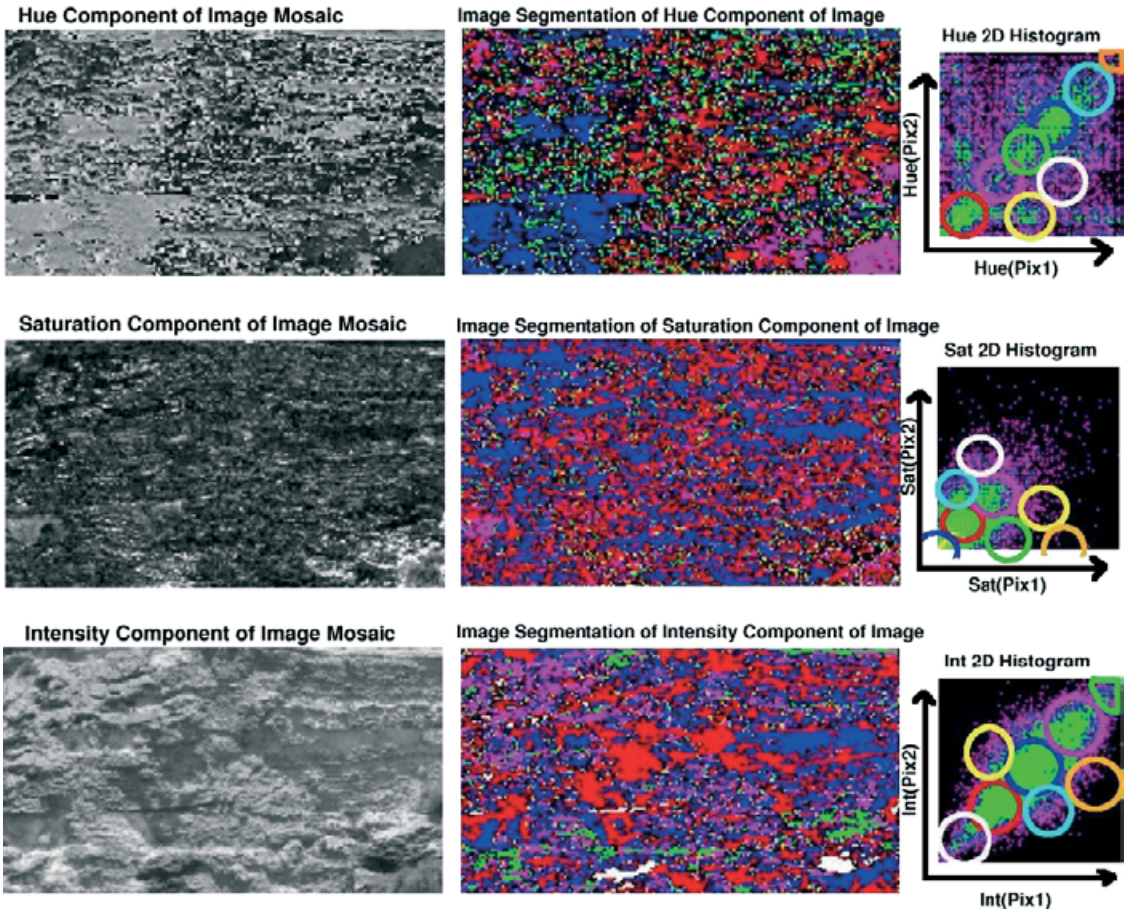

Fig. 9. In the middle column, we show the three image-segmentation maps computed in real-time by the Cyborg Astrobiologist system, based on the original hue, saturation, and intensity (H,S,I) mosaics in the left column, and the derived two-dimensional co-occurrence histograms shown in the right column. The wearable computer made this and all other computations for the original 3×4 mosaic (108×192 pixels, shown in Fig. 7) in about 2 min after the initial acquisition of the mosaic sub-images was completed. The coloured regions in each of the three image-segmentation maps correspond to pixels and their neighbours in that map that have similar statistical properties in their two-point correlation values, as shown by the circles of corresponding colours in the two-dimensional histograms in the column on the right. The red regions in the segmentation maps correspond to the mono-statistical regions with the largest area in this mosaic image; the red regions are the least 'uncommon' pixels in the mosaic. The blue regions correspond to the mono-statistical regions with the second largest area in this mosaic image; the blue regions are the second least 'uncommon' pixels in the mosaic. Similarly for the purple, green, cyan, yellow, white, and orange areas. The pixels in the black regions have failed to be segmented by the segmentation algorithm.

analysis of the dark and wet spot shown in Fig. 8 (a), since the mosaic of this dark and wet spot is less complex than mosaic shown in Fig. 8(b). In retrospect, given the current advancement of the image-analysis software, we should have taken the mosaic from a distance of 3–5 m instead of 10 m, in order to focus *only* on the dark and wet spot, and not on its periphery. With the current image-analysis software, the computer correctly chose to find the periphery interesting because the periphery was the smaller region of the image and because the larger and more common dark and wet spot region was rather homogeneous in its hue and intensity properties. The dark and wet spot had particularly diverse properties of saturation; this information could be useful to develop future interest map algorithms for the Cyborg Astrobiologist.

More advanced software could be developed to better handle the close-up real-time interest-map analysis of such imagery as shown in Figs 8(a) and 8(b). Options to be included in such software development include the following.

- Reintroduce the cortical interest map adaptive weighting algorithm or similar algorithms (i.e. see Chung *et al.* (2004) for a recent discussion) in order to determine the optimal weighting of the H,S,I uncommon maps for this imagery. This way the system could pay more attention to the saturation uncommon map if it deems that this map is the salient map.
- Enhance the interest map so that it includes maps of edges or other image features in addition to the uncommon maps of hue, saturation, and intensity. This could be useful in searching for interesting boundaries or 'contacts' between homogeneous regions, which is a standard technique of field geologists.
- Add hardware and software to the Cyborg Astrobiologist so that it can make intelligent use of its zoom lens. We

https://doi.org/10.1017/S147355040500220X Published online by Cambridge University Press

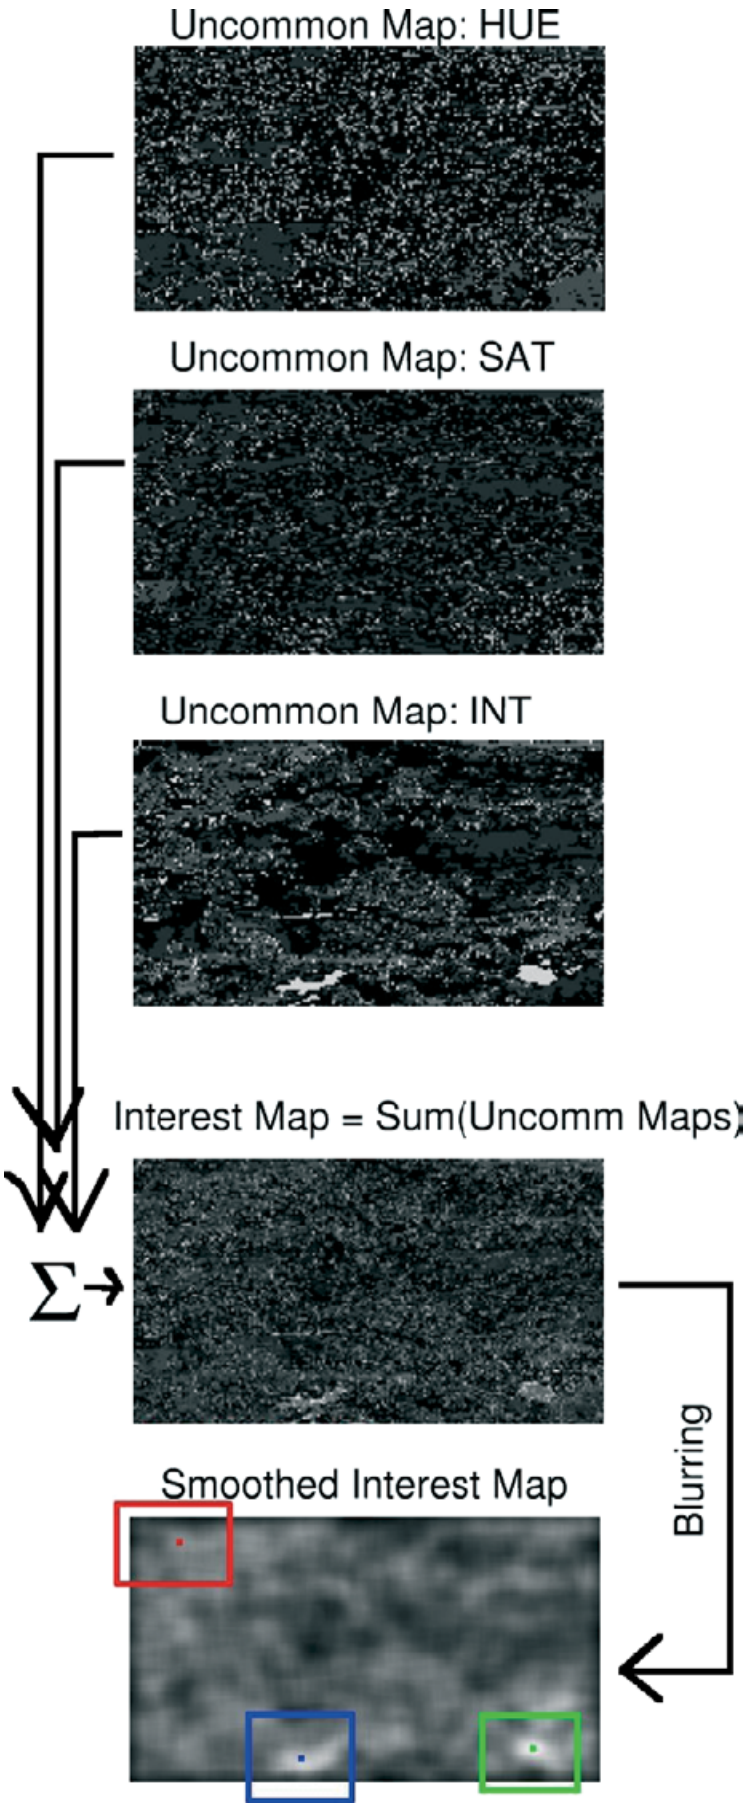

Fig. 10. For legend see opposite page.

https://doi.org/10.1017/S147355040500220X Published online by Cambridge University Press

would need to study and develop the camera's LANC communication interface, or possibly determine whether the zoom lens could be controlled over its Firewire/IEEE1394 communication cable. With accompanying intelligent-zooming software, the system could have corrected the human's mistake in tripod placement and decided to zoom further in, to focus only on the shadowy part of the dark and wet spot (which was determined to be the most interesting point at a distance of 60 m), rather than the periphery of the entire dark and wet spot.

- Enhance the Cyborg Astrobiologist system so that it has a memory of the image segmentations performed at a greater distance or at a lower magnification of the zoom lens. Then, when moving to a closer tripod position or a higher level of zoom-magnification, register the new imagery or the new segmentation maps with the coarser-resolution imagery and segmentation maps. Finally, tell the system to mask out, ignore or de-emphasize those parts of the higher-resolution imagery which were part of the low-interest segments of the coarser, more distant segmentation maps, so that it concentrates on those features that it determined to be interesting at coarse resolution and greater distance.
- Develop true methods of image interpretation, which can define the dark and wet spot region in Fig. 8(a) as one aspect of the image and the periphery of the dark and wet spot as another aspect of the image. Then, determine the relationship between these two aspects. Similarly, for Fig. 8(b), determine, describe, and perhaps explain the relationship between the several different aspects of this image (dark and wet spot; frothy white region to the lower left of the dark and wet spot; green plant to the left of this frothy white area; dead brown plant to the left of this green plant; and dry, unvegetated areas in the rest of the image).

Discussion and conclusions

Both of the geologists on our team concur with the judgement of the Cyborg Astrobiologist software system, that the two dark and wet spots on the cliff wall were the most interesting spots during the second mission. However, the two geologists also state that this largely depends on the aims of study for the geological field trip: if the aim of the study is to search for hydrological features, then these two dark and wet spots are certainly interesting; however, if faulting between layers is of interest, then the dark and wet spots would not retain as much interest. One question which we have thus far left unstudied is 'what would the Cyborg Astrobiologist system have found interesting during the second mission if the two dark and wet spots had not been present?'. It is possible that it would again have found some dark shadow particularly interesting, but with the improvements made to the system between the first and second mission, it is also possible that it could have found a different feature of the cliff wall more interesting.

Thought experiment: blueberry fields in the Meridiani Planum

One of the remarkable discoveries made by the MER team this year is the existence of small, dark and round pebbles (known as 'blueberries') in many places near the landing site of the Opportunity rover, Meridiani Planum. The Opportunity rover was sent to Meridiani Planum because this was the main site on Mars that had been observed by the thermal emission spectrometer on the Mars Global Surveyor orbiter to have a large quantity of grey hematite (Christensen *et al*. 2000). Gray hematite is an iron-bearing mineral which often forms in the presence of water, particularly as a precipitate in cold standing water (Glotch *et al*. 2004). The so-called blueberries that Opportunity discovered at Meridiani Planum were also observed to contain significant hematite levels, based on studies from both of the mini-TES and Mössbauer spectrometer instruments onboard the Opportunity rover (Squyres 2004b). There is significant excitement in the planetary geology community about this discovery, partly due to the existence of a place on Earth (Utah) where similar blueberries have been formed (Catling 2004; Chan *et al*. 2004; Moore 2004; Squyres & Athena Science Team 2004).

Such a geological phenomenon suggests a thought experiment that could be applied to our Cyborg Astrobiologist system as it currently exists: 'if the Cyborg Astrobiologist system was taken to a field covered with Meridiani/Utah blueberries, would it find the blueberries to be interesting?'. The brief, unqualified answer to this thought experiment is 'no'. If the blueberries covered the landing site, so that their coverage exceeded some percentage (assume for argument's sake that the areal coverage of blueberries at the location of the Cyborg Astrobiologist is 80%), then an image mosaic taken by the Cyborg Astrobiologist would be dominated by dark blueberries. Therefore, the Cyborg Astrobiologist's uncommon map algorithm, as it now exists, would likely search the remaining 20% of the area for the smallest colour class of pixels, and it would ignore the blueberries.

However, if the Cyborg Astrobiologist was allowed to approach the field of blueberries from a distance (i.e. from orbit or from a neighbouring area without blueberries), then it is probable that the Cyborg Astrobiologist would at some point find the blueberries to be new and interesting. This is because

Fig. 10. These are the uncommon maps for the mosaic shown in Fig. 7, based on the region sizes determined by the image-segmentation algorithm shown in Fig. 9. Also shown is the interest map, i.e. the unweighted sum of the three uncommon maps. We blur the original interest map before determining the 'most interesting' points. These 'most interesting' points are then sent to the camera's pan-tilt motor in order to acquire and save-to-disk three higher-resolution RGB colour images of the small areas in the image around the interest points (see Fig. 7). Green is the most interesting point, blue is second most interesting, and red is third most interesting.

https://doi.org/10.1017/S147355040500220X Published online by Cambridge University Press

blueberries do not cover all of Mars, and so at some spatial scale, they only have a small areal coverage. So the qualified answer to the thought experiment is 'yes'.

Furthermore, we could program more advanced renditions of the software inside of the Cyborg Astrobiologist so that it could find the blueberries interesting even if the Cyborg Astrobiologist was not allowed to approach the blueberry field from a distance. Such algorithm enhancements could include the texture and edge information suggested above, as well as a possible algorithm to search for uncommon shapes. With spectrometer information (i.e. visible/near-infrared, mid-infrared, Raman, and/or Mössbauer) and with biased searching for particularly interesting spectra, textures, rock shapes, and/or rock juxtapositions, the odds of success could also be maximized, either in hunting for blueberries or in serendipitously finding such blueberries.

These advances in our system could take some years to implement, but we will surely implement some of them in the near-term. In the meantime, as discussed above, the valued concept of a geological approach will remain our solution to the 'blueberry-field' thought experiment. Nevertheless, such a thought experiment does point out a significant inadequacy of our simple computer-vision algorithm which searches for small regions of the colour imagery, as discussed above. Such simple algorithms need to be enhanced and brought into relation with other algorithms, and perhaps with a visual memory of past imagery (i.e. the VAMPIRE project (Wrede *et al*. 2004)).

A visual memory would also be useful to record and track changes in the areal coverage of components in the scenery, as mosaics or images are acquired from different camera positions or for different outcrops. In a different outcrop, the dark wet spots might dominate the whole outcrop, or the whitish colour from the gypsum might not dominate the outcrop. The visual memory could keep track of outcrops studied in the recent past, and inform the user when an uncommon area of a particular image or mosaic is truly novel in areal coverage, when compared with an area of similar colour or texture characteristics in prior imagery.

Outlook

The NEO programming for this Cyborg Astrobiologist/Geologist project was initiated with the SONY Handycam in April 2002. The wearable computer arrived in June 2003, and the HMD arrived in November 2003. We now have a reliably functioning human, hardware, and software Cyborg Geologist system, which is partly robotic with its pan-tilt camera mount. This robotic extension allows the camera to be pointed repeatedly, precisely, and automatically in different directions.

Prior to each of the two field trips to Rivas Vaciamadrid, the Cyborg Geologist system was tested in the grounds of our Spanish research institute (CAB). These tests were mostly on a wide prairie field: by winter, this prairie field was plowed and full of 3–10 cm diameter stones of different, dull colours, and by summer, this prairie field was full of prairie grasses and differently coloured wildflowers. In these nearby field trips, the Cyborg Geologist has provided us with constructive criticism of the basic capabilities of our image-segmentation and interest-map algorithms.

As of our initial field trips to Rivas, the Cyborg Geologist system has accompanied two human field geologists to a site of geological interest (i.e. layered rock outcrops) for more advanced testing of these computer-vision algorithms. We plan several more such trips in the coming months.

Based on the performance of the Cyborg Astrobiologist system during the first mission to Rivas in March 2004 on the outcropping cliffs near Rivas Vaciamadrid, we have decided that the system was paying too much attention to the shadows made by the three-dimensional structure of the cliffs. We hope to improve the Cyborg Astrobiologist system in the next months in order to detect and to pay less attention to shadows. We also hope to upgrade our system to include image-segmentation based on microtexture, and adaptive methods for summing the uncommon maps in order to compute the interest map.

Based on the significantly improved performance of the Cyborg Astrobiologist system during the second mission to Rivas in June 2004, we conclude that the system is now debugged sufficiently so as to be able to produce studies of the utility of particular computer-vision algorithms for geological deployment in the field. We have outlined some possibilities for improvement of the system based on the second field trip, particularly in the improvement in the system-level algorithms needed in order to more intelligently drive the approach of the Cyborg or robotic system towards a complex geological outcrop. These possible system-level improvements include: a better interest-map algorithm, with adaptation and more layers; hardware and software for intelligent use of the camera's zoom lens; a memory of the image segmentation performed at greater distances or lower magnification of the zoom lens; and high-level image-interpretation capabilities.

Now that we have demonstrated that this software and hardware in the Cyborg Astrobiologist system can function for developing and testing computer-vision algorithms for robotic exploration of a geological site, we have some decisions to make as to future directions for this project. Options for these future directions include the following.

- Performing further offline analysis and algorithm development for the imagery obtained at Rivas Vaciamadrid: several of the parameters of the algorithms need testing for their optimality, and further enhancements of the algorithms could be made.
- Optimizing the image-processing and robotic-control code for the current Cyborg Astrobiologist system for speed and memory utilization.
- Further testing of the existing Cyborg geological exploration system at other geological sites with different types of imagery.
- Speeding up the algorithm development by changing the project from being partly a hardware project with cameras and pan-tilt units and fieldwork to being entirely a software project without robotically-obtained image mosaics and

https://doi.org/10.1017/S147355040500220X Published online by Cambridge University Press

without robotic interest-map pointing. With such a change in focus, our algorithms could be significantly enhanced by studying many more types of imagery, for example: from human geologist field studies on the Earth; from robotic geologist field studies on Mars; and from orbiter or flyby studies of our Solar System's moons.

Appendix: software implementation details

Hue, saturation, and intensity are computed in NEO by the standard HSI subroutine, which uses the HSI model described by Foley *et al.* (1990). Hue (H) corresponds roughly to the colour of a pixel, saturation (S) corresponds roughly to the purity of the hue, and intensity (I) corresponds to the total brightness of the pixel. Then each of these three H,S,I image slices is downsampled by a given factor D in each direction (D is typically 2, 4, or 8) in order to conserve memory.

An important but very technical detail in our implementation is that due to the common juxtaposition of pixels of widely different hue (H) in the hue slice (especially for image regions of low saturation), we needed to perform a type of median filter on the Hue slice in non-overlapping $D \times D$ neighbourhoods prior to the downsampling. Such median filtering prior to downsampling was not as important for S and I as it was for H.

The median-filtered and downsampled images are finally combined together into a quasi-mosaic of dimensions $M \times N$ subimages by simply butting the subimages together. We do not perform any mathematical computations to ensure that the edges of the mosaic subimages match well with the edges of the neighbouring subimages.

Following Haralick *et al.* (1973), we developed custom two-dimensional 'co-occurrence' histogramming C code in NST for pairs of pixels of a given separation and orientation, using min/max stretching for preprocessing prior to creating the histograms. For these studies, we chose to histogram pixels separated by 1 pixel width and which are horizontal neighbours, but these two-dimensional histograms are separately made (Fig. 9) for each of the hue, saturation, and intensity image slices. For other studies, we have chosen, for example, to make eight simultaneous and separate two-dimensional histograms for pixels separated by 5 pixels and by 1 pixel, and with pixel orientations of 0°, 45°, 90°, and 135°. Perhaps in the future, we can implement and test something akin to a wavelet analyser of texture.

With the three two-dimensional histograms (for H,S,I), we use the standard *locate_peaks* peak finder in NEO for searching for the eight largest peaks in the two-dimensional histogram. After *locate_peaks* finds a peak, it masks out histogram bins in a disk of a given radius w around that peak, and then it searches the remaining histogram bins for the subsequent peaks of smaller amplitude. The values of w were chosen by trial and error based on previous field work, with the values of w for H and S being 15 stretched units, and the value of w for I being 20 stretched units. This standard *locate_peaks* peak finder may not be very advanced, but it is rather fast and somewhat robust. This robustness is somewhat crucial, since the structure of the two-dimensional Haralick co-occurrence histograms often have structures that might confound a less robust but more advanced peak finder. These two-dimensional histogram structures in the Haralick co-occurrence method, which might confound a more advanced peak finder, include significant variation in the choppiness (or lack of smooth variation) of the co-occurrence histogram as well as significant variation in the width of the true peaks in the co-occurrence histogram. Furthermore, the co-occurrence histograms often only have one true peak, which is somewhat larger in width than the chosen value of w. In this case of only a single true peak, the system will find the first true peak, and it will then find the remaining seven peaks to be the highest points in the histogram that are more distant than w histogram bins from the true peak. This may seem somewhat artificial, but, in practice, it is a decent method to search for those pixels in the image which are truly unusual and interesting, which is our main task in this project. Nonetheless, this peak-finding algorithm does merit significant study and possible revision in future work.

After finding the eight peaks in a two-dimensional histogram of image pixel pairs, we reordered the peaks, ranking the histogram peaks in the order of the number of original-image pixel pairs within their w neighbourhoods (each bin in the two-dimensional histogram counts the number of original-image pixel pairs that have the corresponding values of the ordinate and abscissa). This reordering may seem artificial, but it was necessary given the nature of the *locate_peaks* peak-finding algorithm that we used. To make the possible artificiality clear, the peaks were found based on the largest remaining single pixel values in the two-dimensional co-occurrence histograms, but the peaks were ranked based on the number of pixels in their w neighbourhoods. An alternative is to use a more coherent scheme which finds the eight pre-ranked locations of the w neighbourhoods with largest numbers of pixels. However, we defer this w neighbourhood peak-finding algorithm development for the future since the present system seems to work tolerably well and quickly, and since the artifacts we have observed from using the current scheme of peak finding and peak ordering have not yet suggested the urgency of such an upgraded algorithm.

After the two-dimensional histogramming and peak finding, the images are segmented by our NEO *prog_unit* implementation (custom C code) of the Haddon & Boyce (1990) method. We throw out pixels that have possible conflicts, unlike Haddon and Boyce's more careful work. We handle with some care the pixels on the edges of the subimages in the quasi-mosaic. However, in practice, these edge pixels cause the most problems with the image segmentation when the motor for the pan-tilt unit is taking mosaicking steps that do not match well with the optical zoom factor of the camera, so that the edges of the subimages are easily visible in the quasi-mosaic. A secondary source of subimage edge problems for this image segmentation occurs when there are cumulus clouds in the sky, occasionally occulting the Sun, which causes some of the subimages to have different lighting

https://doi.org/10.1017/S147355040500220X Published online by Cambridge University Press

properties to their neighbours. If a pixel and its horizontal neighbour have values that sets it in one peak's neighbourhood (i.e. peak number s of the eight ranked peaks of that two-dimensional histogram), then the pixel is said to be in segment s of that image. The segmentation maps for that image are therefore coloured by eight different colours, with red for the 'most common' segment $s=1$ (which has the most pixels in it), blue for second 'most common' segment $s=2$, etc. See Fig. 9 for an example.

Acknowledgements

P. McGuire, J. Ormö and E. Díaz Martínez would all like to thank the Ramon y Cajal Fellowship program in Spain, as well as certain individuals for assistance or conversations: Virginia Souza-Egipsy, Eduardo Sebastián Martínez, María Paz Zorzano Mier, Carmen Córdoba Jabonero, Kai Neuffer, Antonino Giaquinta, Fernando Camps Martínez, Alain Lepinette Malvitte, Josefina Torres Redondo, Víctor R. Ruiz, Julio José Romeral Planelló, Gemma Delicado, Jesús Martínez Frías, Irene Schneider, Gloria Gallego, Carmen González, Ramon Fernández, Coronel Angel Santamaria, Carol Stoker, Paula Grunthaner, Maxwell D. Walter, Fernando Ayllón Quevedo, Javier Martín Soler, Juan Pérez Mercader, Jörg Walter, Claudia Noelker, Gunther Heidemann, Robert Haschke, Robert Rae, Jonathan Lunine, and two anonymous referees. The work of J. Ormö was partially supported by a grant from the Spanish Ministry for Science and Technology (AYA2003-01203). The equipment used in this work was purchased by grants to our Center for Astrobiology from its sponsoring research organizations, CSIC and INTA.

References

Apostolopoulos, D., Wagner, M.D., Shamah, B., Pedersen, L., Shillcutt, K. & Whittaker, W.L. (2000). Technology and field demonstration of robotic search for Antarctic meteorites. *Int. J. Robot. Res.* **19(11)**, 1015–1032.

Arvidson, R.E. *et al.* (2003). Physical properties and localization investigations associated with the 2003 Mars Exploration rovers. *J. Geophys. Res.* (*Planets*) **108**, 11–1.

Arvidson, R.E. *et al.* (2004). Localization and physical properties experiments conducted by Spirit at Gusev Crater. *Science* **305**, 821–824.

Batavia, P.H. & Singh, S. (2001). Obstacle detection using adaptive color segmentation and color stereo homography. In *Proc. IEEE Conf. on Robotics and Automation*, Seoul, Korea, pp. 705–710.

Batlle, J., Casals, A., Freixenet, J. & Martí, J. (2000). A review for recognizing natural objects in colour images of outdoor scenes. *Image Vision Comput.* **18**, 515–530.

Bell, J.F. *et al.* (2003). Mars Exploration Rover Athena Panoramic Camera (Pancam) investigation. *J. Geophys. Res.* (*Planets*) **108**, 4–1.

Bell, J.F. *et al.* (2004). Pancam multispectral imaging results from the Spirit Rover at Gusev Crater. *Science* **305**, 800–807.

Cabrol, N.A. *et al.* (2001). Nomad Rover field experiment, Atacama Desert, Chile 2. Identification of paleolife evidence using a robotic vehicle: lessons and recommendations for a Mars sample return mission. *J. Geophys. Res.* **106** E4, 7807–7811.

Calvo, J.P., Alonso, A.M. & Garcia del Cura, A.M. (1989). Models of Miocene marginal lacustrine sedimentation in response to varied depositional regimes and source areas in the Madrid Basin (central Spain). *Paleogeogr. Paleoclim. Paleoecol.* **70**, 199–214.

Castilla Cañamero, G. (2001). Informe sobre las prácticas profesionales realizadas en el Centro de Educación Ambiental: El Campillo. Internal report to the Consejería de Medio Ambiente de la Comunidad de Madrid.

Catling, D.C. (2004). Planetary science: on Earth, as it is on Mars? *Nature* **429**, 707–708.

Chan, M.A., Beitler, B., Parry, W.T., Ormö, J. & Komatsu, G. (2004). A possible terrestrial analogue for haematite concretions on Mars. *Nature* **429**, 731–734.

Cheeseman, P., Kelly, J., Self, M., Stutz, J., Taylor, W. & Freeman, D. (1988). AutoClass: a Bayesian classification system. In *Proc. 5th Int. Conf. on Machine Learning*, Ann Arbor, MI, 12–14 June 1988, pp. 54–64. Morgan Kaufmann, San Francisco.

Cheeseman, P. & Stutz, J. (1996). Bayesian classification (AutoClass): theory and results. In *Advances in Knowledge Discovery and Data Mining*, eds Fayyad, U.M., Piatetsky-Shapiro, G., Smyth, P. & Uthurusamy, R., AAAI Press/MIT Press, Cambridge, MA.

Christensen, P.R. *et al.* (2000). Detection of crystalline hematite mineralization on Mars by the Thermal Emission Spectrometer: evidence for near-surface water. *J. Geophys. Res.* **105**, 9623–9642.

Christensen, P.R. *et al.* (2004). Initial results from the mini-TES experiment in Gusev Crater from the Spirit Rover. *Science* **305**, 837–842.

Chung, A.J., Deligianni, F., Hu, X.P. & Yang, G.Z. (2004). Visual feature extraction via eye tracking for saliency driven 2D/3D registration. In *Conf. on Eye Tracking Research and Applications* (*ETRA04*), San Antonio, TX, pp. 49–54.

Corsetti, F.A. & Storrie-Lombardi, M.C. (2003). Lossless compression of stromatolite images: a biogenicity index? *Astrobiology* **3(4)**, 649–655.

Crawford, J. (2002). *NASA Applications of Autonomy Technology*, NASA Ames Research Center, Intelligent Data Understanding Seminar. http://is.arc.nasa.gov/IDU/slides/Crawford_Aut02c.pdf.

Crawford, J. & Tamppari, L.K. (2002). *Mars Science Laboratory – Autonomy Requirements Analysis* NASA Ames Research Center, Intelligent Data Understanding Seminar. http://is.arc.nasa.gov/IDU/slides/Crawford_MSL02c.pdf.

Crisp, J.A., Adler, M., Matijevic, J.R., Squyres, S.W., Arvidson, R.E. & Kass, D.M. (2003). Mars Exploration Rover mission. *J. Geophys. Res.* (*Planets*) **108**, 2–1.

Crumpler, L. *et al.* (2004). MER field geologic traverse in Gusev Crater, Mars: initial results from the perspective of Spirit. *Lunar Planet. Inst. Conf. Abstracts* **35**, 2183.

Foley, J., van Dam, A., Fiener, S. & Hughes, J. (1990). *Computer Graphics: Principles and Practice*, 2nd edn. Addison-Wesley, Reading, MA.

Förstner, W. (1986). A feature based algorithm for image matching. *Int. Archives Photogram. Remote Sensing* **26(3)**, 150–166.

Förstner, W. & Gülch, E. (1987). A fast operator for detection and precise location of distinct points, corners and centres of circular features. In *Proc. Intercommission Conf. on Fast Processing of Photogrammetric Data*, Interlaken, Switzerland, pp. 281–305.

Freixenet, J., Muñoz, X., Martí, J. & Lladó, X. (2004). Color texture segmentation by region-boundary cooperation. In *ECCV 2004, Proc. 8th Eur. Conf. on Computer Vision*, Prague, Czech Republic, Lecture Notes in Computer Science vol. 3022, eds Pajdla, T. & Matas, J., pp. 250–261. Springer, Berlin. Also available at: http://homepages.inf.ed.ac.uk/rbf/CVonline/LOCAL_COPIES/FREIXENET1/eccv04.html.

Glotch, T.D. *et al.* (2004). Hematite at Meridiani Planum: detailed spectroscopic observations and testable hypotheses. *Lunar Planet. Inst. Conf. Abstracts* **35**, 2168.

Goldberg, S.B., Maimone, M.W. & Matthies, L. (2002). Stereo vision and rover navigation software for planetary exploration. *Proc. 2002 IEEE Aerospace Confer. Proc.*, Big Sky, Montana, vol. 5, pp. 2025–2036.

Golombek, M. *et al.* (2004). Preliminary assessment of Mars Exploration Rover landing site predictions. *Lunar Planet. Inst. Conf. Abstracts* **35**, 2185.

Greeley, R. *et al.* (2004). Wind-related processes detected by the Spirit Rover at Gusev Crater, Mars. *Science* **305**, 810–821.

Gulick, V.C., Morris, R.L., Ruzon, M.A. & Roush, T.L. (2001). Autonomous image analyses during the 1999 Marsokhod rover field test. *J. Geophys. Res.* **106**, 7745–7764.

Gulick, V.C., Morris, R.L., Gazis, P.R., Bishop, J.L. & Alena, R.L. (2002a). The Geologist's Field Assistant: developing an innovative science analysis system for exploring the surface of Mars. *AGU Fall Meeting Abstracts*, **A364**.

Gulick, V.C., Morris, R.L., Bishop, J., Gazis, P., Alena, R. & Sierhuis, M. (2002b). Geologist's Field Assistant: developing image and spectral analyses algorithms for remote science exploration. *Lunar Planet. Inst. Conf. Abstracts* **33**, 1961.

Gulick, V.C., Morris, R.L., Gazis, P., Bishop, J.L., Alena, R., Hart, S.D. & Horton, A. (2003). Automated rock identification for future Mars exploration missions. *Lunar Planet. Inst. Conf. Abstracts* **34**, 2103.

Gulick, V.C., Hart, S.D., Shi, X. & Siegel, V.L. (2004). Developing an automated science analysis system for Mars surface exploration for MSL and beyond. *Lunar Planet. Inst. Conf. Abstracts* **35**, 2121.

Haddon, J.F. & Boyce, J.F. (1990). Image segmentation by unifying region and boundary information. *IEEE Trans. Pattern Anal. Mach. Intell.* **12**(**10**), 929–948.

Haralick, R.M., Shanmugan, K. & Dinstein, I. (1973). Texture features for image classification. *IEEE Trans. Syst., Man., Cybernet.* **3**(**6**), 610–621.

Heninger, B., Sander, M., Simmonds, J., Palluconi, F., Muirhead, B. & Whetsel, C. (2003). *Mars Science Laboratory Mission 2009 Landed Science Payload Proposal Information Package* (draft version). http://centauri.larc.nasa.gov/msl/PIP-Drft/_FBO-RevA-031113.pdf.

Herkenhoff, K.E. *et al.* (2004). Textures of the soils and rocks at Gusev Crater from Spirit's microscopic imager. *Science* **305**, 824–827.

Huntsberger, T., Aghazarian, H., Cheng, Y., Baumgartner, E.T., Tunstel, E., Leger, C., Trebi-Ollennu, A. & Schenker, P.S. (2002). Rover autonomy for long range navigation and science data acquisition on planetary surfaces. In *Proc. 2002 IEEE Int. Conf. on Robotics and Automation*, Washington, DC.

IGME (Instituto Geológico y Minero de España). (1975). *Mapa geologico de Españ̃a E 1:50,000*, Arganda (segunda serie, primera edicion), Hoja numero 583, Memoria explicativa, pp. 3–25.

Lee, S. *et al.* (2002). *Spiderman: The Movie*. Marvel Comic Books and Sony Pictures.

Maki, J.N. *et al.* (2003). Mars Exploration Rover engineering cameras. *J. Geophys. Res.* (*Planets*) **108**, 12–1.

McEwen, A. *et al.* (2003). MRO's High Resolution Imaging Science Experiment (HiRISE): science expectations. In *Proc. 6th Int. Conf. on Mars*, p. 3217.

McGreevy, M.W. (1992). The presence of field geologists in Mars-like terrain. *Presence*, **1**(**4**), 375–403.

McGreevy, M.W. (1994). An ethnographic object-oriented analysis of explorer presence in a volcanic terrain environment. *NASA Technical Memorandum #108823*, Ames Research Center, Moffett Field, CA.

McGuire, P.C. *et al.* (2004). Cyborg systems as platforms for computer-vision algorithm-development for astrobiology. In *Proc. 3rd Eur. Workshop on Exo/Astrobiology* (ESA SP-545), Centro de Astrobiologia, Madrid, pp. 141–144. Available at: http://arxiv.org/abs/cs.CV/0401004.

McSween, H.Y. *et al.* (2004). Basaltic rocks analyzed by the Spirit Rover in Gusev Crater. *Science* **305**, 842–845.

Moore, J.M. (2004). Blueberry fields for ever. *Nature* **428**, 711–712.

Nesnas, I., Maimone, M. & Das, H. (1999). Autonomous vision-based manipulation from a rover platform. In *Proc. CIRA Conf.*, Monterey, CA.

Olson, C.F., Matthies, L.H., Schoppers, M. & Maimone, M.W. (2003). Rover navigation using stereo ego-motion. *Robot. Autonomous Syst.* **43**(**4**), 215–229.

Pedersen, L. (2001). Autonomous characterization of unknown environments. In *Proc. 2001 IEEE Int. Conf. on Robotics and Automation*, vol. 1, pp. 277–284.

Rae, R., Fislage, M. & Ritter, H. (1999). Visuelle Aufmerksamkeitssteuerung zur Unterstützung gestikbasierter Mensch–Maschine Interaktion. *KI – Künstliche Intelligenz, Themenheft Aktive Sehsysteme* **1**, 18–24.

Ritter, H. *et al.* (1992, 2002). *The Graphical Simulation Tookit, Neo/NST*. http://www.TechFak.Uni-Bielefeld.DE/techfak/ags/ni/.

Sebe, N., Tian, Q., Loupias, E., Lew, M. & Huang, T.S. (2003). Evaluation of salient points techniques. *Image Vision Comput.* (*Special Issue on Machine Vision*) **21**, 1087–1095.

Squyres, S.W. (2004a). Initial results from the MER Athena science investigation at Gusev Crater and Meridiani Planum. *AGU Spring Meeting Abstracts*, A1.

Squyres, S.W. (2004b). *Lecture Notes from Astrobiology Summer School Course on the MER Rovers*, Universidad Internacional Menendez Pelayo Santander, Spain.

Squyres, S.W. & Athena Science Team (2004). Science results from the MER Athena science investigation at Gusev Crater and Meridiani Planum. Abstract to the Klein Lecture at Astrobiology Science Conference. *Int. J. Astrobiol.* **3** (Suppl.), 5.

Squyres, S.W. *et al.* (2003). Athena Mars rover science investigation. *J. Geophys. Res.* (*Planets*) **108**, 3–1.

Squyres, S.W. *et al.* (2004a). The Spirit Rover's Athena science investigation at Gusev Crater, Mars. *Science* **305**, 794–800.

Squyres, S.W. (2004b). Lecture Notes from the Planet Mars Astrobiology Summer School, Covrs on the MER rovers.

Storrie-Lombardi, M.C., Grigolini, P., Galatolo, S., Tinetti, G., Ignaccolo, M., Allegrini, P. & Corsetti, F.A. (2002). Advanced techniques in complexity analysis for the detection of biosignatures in ancient and modern stromatolites. *Astrobiology* **2**(**4**), 630–631.

Vandapel, N., Chatila, R., Moorehead, S., Lacroix, S., Apostolopoulos, D. & Whittaker, W.L. (2000). Evaluation of computer vision algorithms for autonomous navigation in polar terrains. In *Proc. Int. Conf. on Intelligent Autonomous Systems*.

Volpe, R. (2003). Rover functional autonomy development for the Mars Mobile Science Laboratory. In *Proc. 2003 IEEE Aerospace Conf.*, Big Sky, Montana.

Volpe, R. (2003). *NASA Jet Propulsion Laboratory* (*JPL*) *CLARAty system*.

Wagner, M.D. (2000). Experimenter's Notebook: Robotic Search for Antarctic Meteorites 2000 Expedition. *Technical Report CMU-RI-TR-00-13*, Robotics Institute, Carnegie Mellon University.

Wettergreen, D., Bapna, D., Maimone, M. & Thomas, H. (1999). Developing robotic exploration of the Atacama Desert. *Robot. Autonomous Syst.* **26**(**2–3**), 127–148.

Whittaker, W.L., Bapna, D., Maimone, M. & Rollins, E. (1997). Atacama Desert Trek: a planetary analog field experiment. In *Proc. 4th Int. Symp. on Artificial Intelligence, Robotics and Automation for Space* (*i-SAIRAS'97*), Tokyo, Japan.

Wrede, S., Hanheide, M., Bauckhage, C. & Sagerer, G. (2004). An active memory as a model for information fusion. In *Proc. 7th Int. FUSION Conf.*