https://ntrs.nasa.gov/search.jsp?R=20050157066 2017-09-22T03:20:38+00:00Z

The Ames MER Microscopic Imager Toolkit

Randy Sargent, Matthew Deans, Clayton Kunz
NASA Ames Research Center / QSS Group, Inc.
M/S 269-3
Moffett Field, CA 94035
{rsargent,mdeans,ckunz}@arc.nasa.gov

Michael Sims
NASA Ames Research Center
M/S 269-3
Moffett Field, CA 94035
msims@arc.nasa.gov

Ken Herkenhoff
U.S. Geological Survey
2255 N. Gemini Dr.
Flagstaff, AZ, 86001
kherkenhoff@usgs.gov

Abstract—[1,2]The Mars Exploration Rovers, Spirit and Opportunity, have spent several successful months on Mars, returning gigabytes of images and spectral data to scientists on Earth. One of the instruments on the MER rovers, the Athena Microscopic Imager (MI), is a fixed focus, megapixel camera providing a ±3mm depth of field and a 31x31mm field of view at a working distance of 63 mm from the lens to the object being imaged. In order to maximize the science return from this instrument, we developed the Ames MI Toolkit and supported its use during the primary mission. The MI Toolkit is a set of programs that operate on collections of MI images, with the goal of making the data more understandable to the scientists on the ground. Because of the limited depth of field of the camera, and the often highly variable topography of the terrain being imaged, MI images of a given rock are often taken as a stack, with the Instrument Deployment Device (IDD) moving along a computed normal vector, pausing every few millimeters for the MI to acquire an image. The MI Toolkit provides image registration and focal section merging, which combine these images to form a single, maximally in-focus image, while compensating for changes in lighting as well as parallax due to the motion of the camera. The MI Toolkit also provides a 3-D reconstruction of the surface being imaged using stereo and can embed 2-D MI images as texture maps into 3-D meshes produced by other imagers on board the rover to provide context. The 2-D images and 3-D meshes output from the Toolkit are easily viewed by scientists using other mission tools, such as Viz or the MI Browser.

This paper describes the MI Toolkit in detail, as well as our experience using it with scientists at JPL during the primary MER mission.

TABLE OF CONTENTS

1. INTRODUCTION 1
2. 2-D IMAGE PROCESSING 2
3. 3-D IMAGE PROCESSING 7
4. MICROSCOPIC IMAGE BROWSER 11
5. TOOLKIT USAGE DURING THE MER MISSION ... 13
6.ACKNOWLEDGEMENTS 14
REFERENCES 14
BIOGRAPHY 15

1. INTRODUCTION

The two Mars Exploration Rovers, Sprit and Opportunity, have been operating successfully on Mars for several months, returning gigabytes of data that have enabled scientists to learn a great deal about the history of the surface of the planet. The two rovers carry an identical payload of scientific instruments (the Athena science payload), which includes several fixed-focus cameras, spectrometers, and a rock abrasion tool (RAT). The rovers each have a five degree-of-freedom arm, called the instrument deployment device (IDD), at the end of which several of the instruments are mounted, including the microscopic imager (MI) [7]. The MI has a 20mm focal length lens at f/15, providing a +/- 3mm depth-of-field, and can view an area of approximately 31mm square at the working distance of 63 mm from the lens to the target object. MI images have a resolution of 1024x1024 pixels, giving a resolution of approximately 3 microns per pixel.

The MI functions analogously to a geologist's hand lens, allowing scientists on Earth to examine the coarse crystalline structures of the surfaces of rocks, as well as properties revealed after abrasion by the RAT. The MI is also used as a diagnostic tool, to examine the magnets of the

[1] 0-7803-8870-4/05/$20.00© 2005 IEEE
[2] IEEEAC paper #1227, Version 13, Updated December 9, 2004.

rovers' Mössbauer spectrometer.

The MI is often used to image areas larger than its working field of view, and objects that have more than 3mm of relief; in these cases the IDD is moved around while the MI takes pictures, returning either a mosaic of images of a large field of view (FOV) scene, or a set of images of the same scene taken at different distances from it, which we refer to as a *stack*. The volume of data taken by the MI can therefore accumulate rapidly, challenging both the available communications bandwidth for downlink, and the amount of time scientists on the ground have to understand what the rovers are seeing in time to make planning decisions for the next day. We developed the Ames MI Toolkit at NASA Ames Research Center to help scientists understand the images taken by the MI, and to help planners make decisions about how many images should be taken, and the amount of lossy image compression that should be applied to ensure maximum useful information return.

The Ames MI Toolkit is a set of image processing tools for combining and organizing multiple images into coherent and easily understood models. The toolkit provides image registration and focal section merging for combining a stack of images taken at different target distances into a single, globally in-focus image. It also provides tools for 3-D reconstruction of a surface, using stereo. The toolkit does not provide 2-D mosaicking or colorization of images; USGS and Cornell worked on these capabilities.

We will discuss the 2-D and 3-D uses of the MI Toolkit in turn, including the relationship between the components of the toolkit and the state of the art in computer vision. We will follow with a discussion of how the toolkit was put to use during the mission.

2. 2-D IMAGE PROCESSING

The Ames MI Toolkit automatically performs image registration and *focal section merging*, the combining several images of a surface into a single maximally in-focus image. The toolkit also can analyze images to find shadows and attempt to adjust the image to appear as if the shadow is removed.

Image Registration

Focal section merging requires that we first find corresponding pixels in a stack of MI images. For the best focal section merging, we want a dense, subpixel estimate for correspondences, which can be expensive to compute when image motion is large and there are no *a priori* epipolar constraints. In order to facilitate the search for these correspondences, we first register the images up to a homography. If the scene is relatively planar, this homography accounts for most of the image plane motion of scene points, and the job of the pixel by pixel correspondence search is made much simpler.

The image registration uses a nonlinear optimization method to recover the 8 parameters describing the homography. The homography H describes the relationship between pixel coordinates in images I_1 and I_2 as

$$I_1(x) = I_2(x') = I_2(Hx) \tag{1}$$

where x and x' are projective coordinates in 2D, and the equality is up to a scale factor. The matrix H describing the homography has the form

$$H = \begin{bmatrix} a & b & c \\ d & e & f \\ g & h & 1 \end{bmatrix} \tag{2}$$

We choose the lower right entry to be unity because the homography is only defined up to a scale factor, and search for the other 8 parameters. In addition, the two images may have slightly different exposures, which we account for using a linear function of the pixel values, and each pixel will contain some random noise, so that equation (1) becomes

$$I_1(x) = \alpha I_2(Hx) + \beta + \varepsilon_x \tag{3}$$

For convenience, the 8 parameters of the homography and the linear coefficients α and β can be collected into the parameter vector $\theta=(a,b,c,d,e,f,\alpha,\beta)$. We now define a least squares cost function over θ,

$$J(\theta) = \sum_x \left[I_1(x) - \alpha I_2(Hx) + \beta\right]^2 \tag{4}$$

and minimize $J(\theta)$ using Levenberg-Marquardt.

Working with full resolution images presents two problems. First, the cost function in (4) may have lots of local minima which can trap the minimization. Second, the derivatives and cost function can take a long time to compute. For both of these reasons, we use a coarse-to-fine approach based on image pyramids.

The image pyramid is constructed for both I_1 and I_2 up to level 3, or subsampling to 2^{-3} the original image dimensions. These subsampled images are aligned only up to rotation and scale, i.e.

$$H_3 = \begin{bmatrix} \cos(\phi) & -\sin(\phi) & u \\ \sin(\phi) & \cos(\phi) & v \\ 0 & 0 & 1 \end{bmatrix} \tag{5}$$

The level 3 registration starts with H_3 initialized to the identity matrix. Once the level 3 registration converges, the level 2 images are registered, again up to rotation and scale. In order to initialize the search at level 2, we start with the

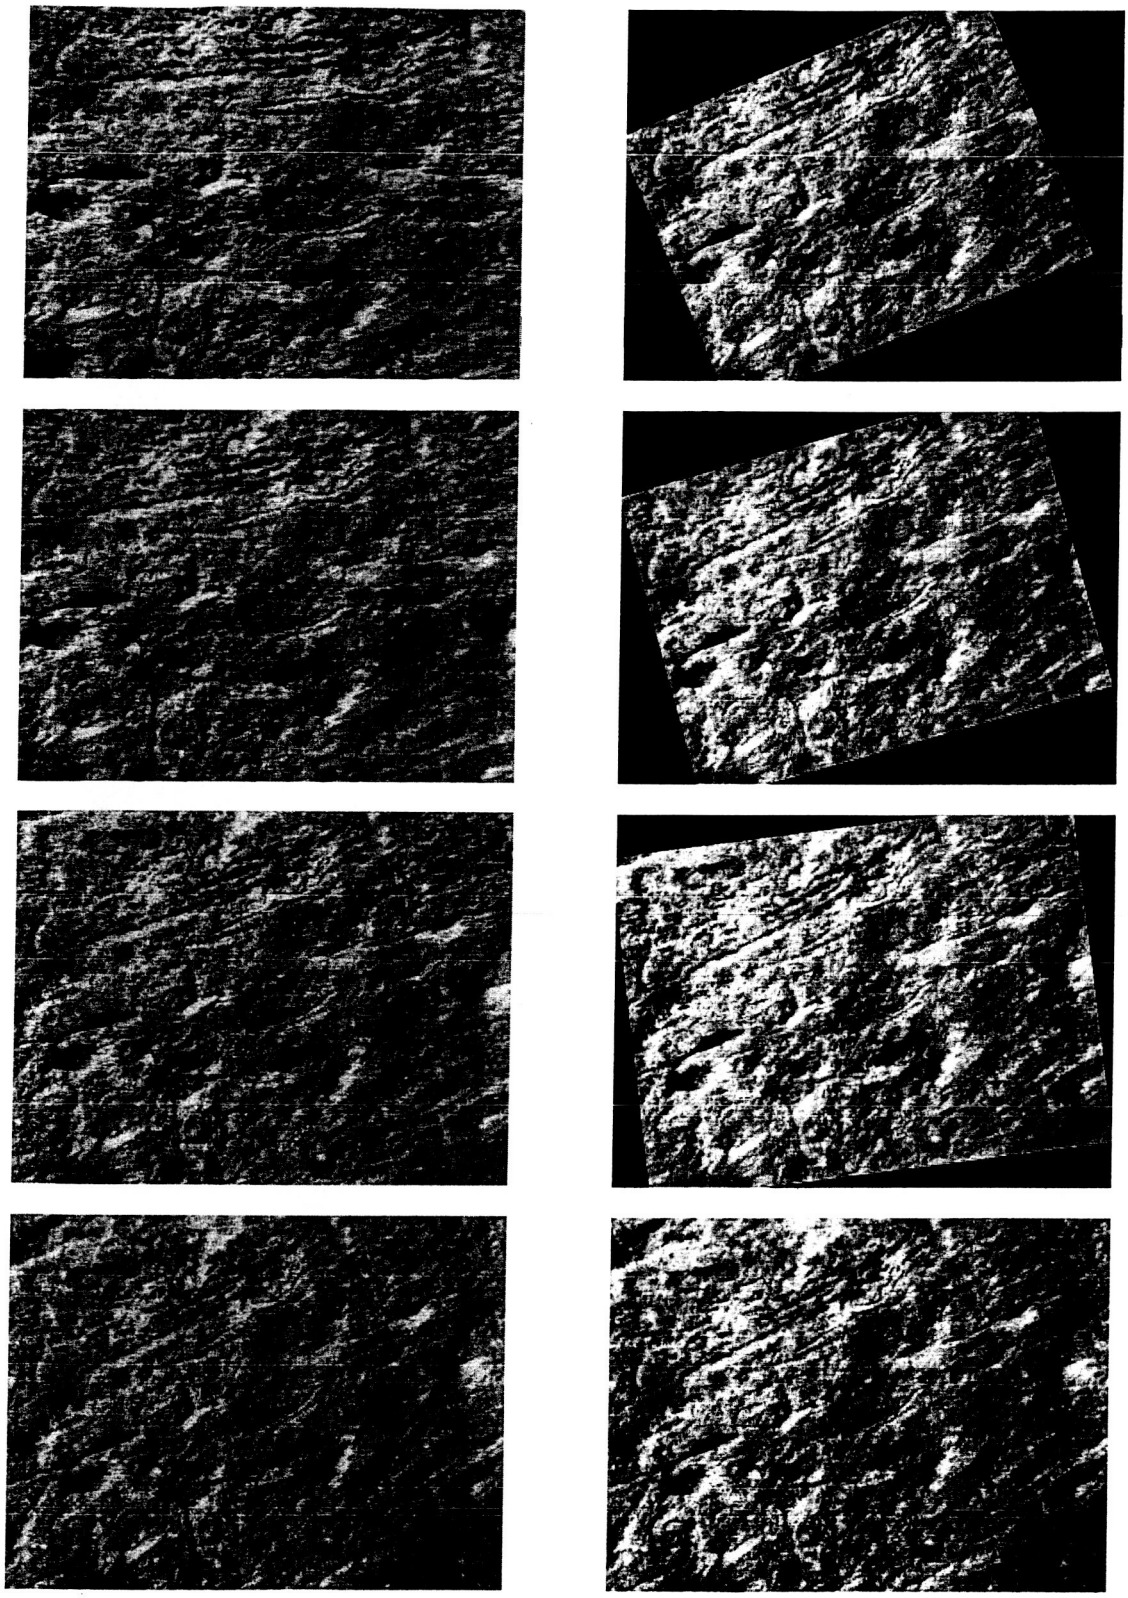

Figure 1 Example MI image sequence with the camera rotating and moving away from a rock surface. Images on left are original images. Images on right are the corresponding image registration results. Note that the images from closer to the surface have missing data when aligned with the image from farther out, shown here as black borders. *Original images courtesy of Cornell University*

rotation angle recovered at level 3 and multiply the translation parameters by 2 to account for the difference in scale at the next pyramid level. The same procedure is followed at level 1, and the final registration is a search for the full homography at level 0 (full resolution), initialized by H_1.

Figure **1** shows an example image sequence using an engineering model of the MER MI taken at Cornell University, as well as the result of the registration.

Dense Correspondence

A stack of images registered using homography is a useful product. Among other things it allows scientists to browse a registered stack, scrolling from near to far focus, without the distraction of large image motions. However, homography does not correct for parallax due to camera motion relative to a non-planar surface.

Although parallax is typically associated with lateral camera motion, it is also present for motion along the camera pointing direction, increasing towards the edges of the image. Parallax motions of several pixels are commonly observed between successive MI images, while even a fractional pixel is enough to cause artifacts when performing the focal section merge.

We initially used a simple window correlation search to find correspondences between images. For a given pixel in the first image, a pixel is chosen in the second image minimizing the sum square of differences of corresponding pixels in a window around each pixel.

Correcting each image in a stack in this way, we created a set of images we hoped would exhibit no relative motion when viewed sequentially, only changes in focus. Disappointingly, this wasn't the case. Small motions (less than a pixel) were apparent when animating between images, and this miscorrelation caused features to grow or shrink in the resulting focal section merge.

We first modified the window correlation search to calculate a subpixel match location by modeling the neighborhood of the optimal match as a quadratic basin and finding its minimum. This improved the results, but unfortunately small motions persisted in areas of the image around large depth relief.

The continued problem turned out to be the parallax itself; different portions of the support window around a pixel were moving by different amounts. Depending on the texture of the scene, the area very close to the pixel might dominate the match solution, or an area towards the edge of the window, which moved differently. Minimizing the support window size reduced this effect at the cost of spurious incorrect solutions and noise in the quadratic subpixel solution.

To solve our problem, we developed a new technique for maintaining support window size while removing the problem of parallax. We used an iterative approach to solving the correspondences. For the first iteration, a complete correlation search is performed between the two images, which will be called *A* and *B*. This correlation provides our initial estimate for parallax motion. For each successive iteration, we use the parallax estimation from the previous iteration to warp image *A* into image *A'*, and then perform a correlation between *A'* and B. This correlation provides a correction to our current estimate of parallax

Figure 2 Two original images and the magnitude of the parallax correction applied (black=0 pixels, white=23 pixels)

motion.

As the iterations continue, A' approaches B in appearance, and as it does, the error in correlation minimum due to parallax reduces. The resulting solution produces images without perceptible subpixel motion, allowing us to generate focal section merges without perceptible feature size changes. Figure 2 shows the magnitude of the parallax correction at each pixel after homography.

Shadow Correction

Due to operational constraints and time constraints on some science objectives, the MER MI is occasionally used in situations where the science target is shadowed by the rover or the instrument itself. This can present a problem for image registration and sharpness value computation since the same scene region might be imaged in direct illumination in one image and in cast shadow in subsequent images. The robust estimator used in image registration reduces the effect on alignment but this only partially addresses the issue for creating a single complete best view. In order to reduce the effect of shadows, we developed a shadow estimation and removal tool.

We begin by assuming that each pixel in the image is either directly illuminated or in shadow, so the estimation of shadow/illuminated regions is a binary labeling problem as in [6]. Furthermore, we assume that each region is characterized by a statistical distribution of pixel intensity values which can be reasonably represented using a Gaussian distribution. The particular model does not need to match well, just well enough to aid in the binary labeling problem. Lastly we assume that regions are smooth, that is we assume an a priori probability that neighboring pixels share labels, implying that neighboring pixels tend to be simultaneously in shadow or in illumination.

Combining the *a priori* smoothness with the label likelihood leads to the posterior distribution over labels.

$$P(f \mid I) = P_{data}(I \mid f) P_{smooth}(f)$$

(5)

$$= \prod_i e^{-\frac{(i-\mu_f)^2}{2\sigma^2}} \prod_{ij} p(f_i, f_j) \qquad (6)$$

where the first product is the likelihood of label f for the observed pixel intensity, and the second product is the probability that two neighboring pixels share the same label. Maximizing (5) is the same as minimizing the negative log, so we take the log to turn the product into a sum and to turn the exponential into a simple energy function with quadratic and constant terms, yielding

$$E(f) = E_{data}(f) + E_{smooth}(f) \qquad (7)$$

with

$$E_{data}(f) = \sum_i \frac{(i-\mu_f)^2}{2\sigma_f^2} \qquad (8)$$

and

$$E_{smooth}(f) = -\sum_{ij} \log\left(p(f_i, f_j)\right) \qquad (9)$$

Equation (7) can be solved in closed form since the labeling f is binary [6][2].

The shadow correction requires three steps. The first step is to recover the distributions of pixel intensities in illumination and in shadow. We compute the histogram of image pixel intensities. This histogram is fit to a mixture model using EM to recover the mean and variance of the two Gaussians [1]. Figure 3 shows an example histogram and the recovered mixture model with shadowed regions modeled with the normal distribution given by *N(80.19,18.53)* and illuminated regions modeled with *N(157.28,20.04)*. Note that in this example the mixture model fits the image histogram well, but in practice the mixture model only needs to match well enough to aid in the binary label problem.

The second step is to solve equation (6) for the labeling. The values of $E_{data}(f)$ and + $E_{smooth}(f)$ are computed and used to construct a network flow problem whose solution is the labeling that minimizes (6). Our implementation uses Boykov's public domain network flow solver [2].

Figure 4 shows an example MI image with cast shadows. The histogram of pixel values and mixture model recovered by EM are shown in Figure 3. A maximum likelihood labeling (ignoring the smoothness heuristic) is shown in Figure 5, where many mislabeled pixels appear as isolated points for which the wrong label has higher likelihood than the correct label. Figure 6 shows the result of applying the graph cut to minimize (6). Finally, Figure 7 shows the result of applying a linear photometric correction to each pixel labeled as shadowed.

Our shadow labelling and compensation method uses a simplistic model of image formation in the presence of cast shadows, and the result leads to strong artifacts on shadow boundaries, as seen in Figure 7. In reality, the pixels near the edge of the shadowed region are not as strongly affected by the object casting the shadow, and the compensation should be less. This is an area for future work.

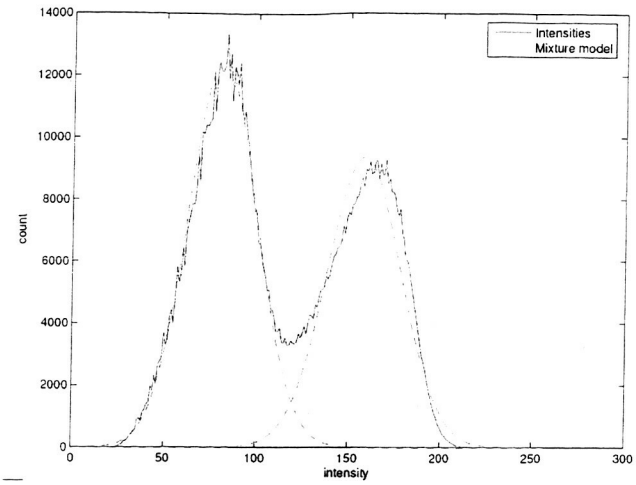

Figure 3 Pixel histogram for image in Figure 4

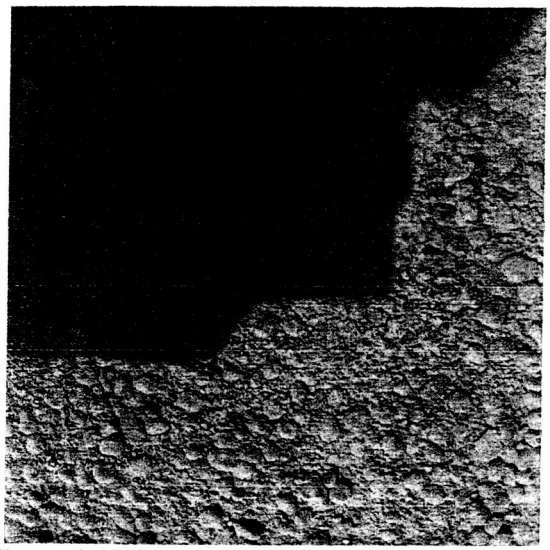

Figure 4 MI image with direct illumination and cast shadows.

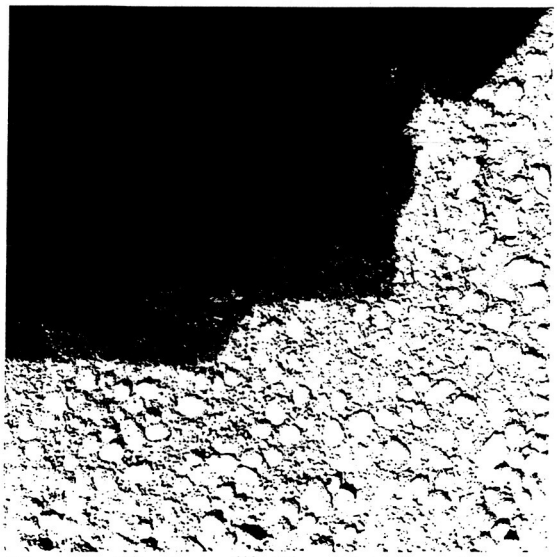

Figure 5 Maximum likelihood labeling of pixels.

Figure 6 Maximum a posteriori labeling using graph cut.

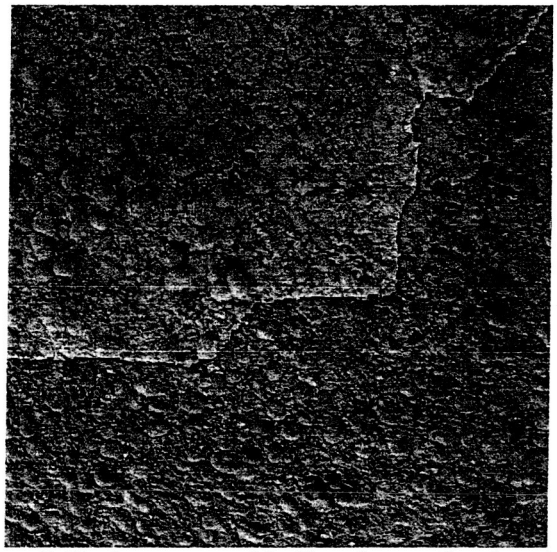

Figure 7 Shadow corrected by linearly transforming pixels labeled as being in shadow.

Focal Section Merging

The original motivation for the development of the MI toolkit is the problem of focal section merging, which is necessary because of the limited depth-of-field of the MI's optics. While focus is a good cue for relative depth for humans as they flip back and forth between images in a vertical sequence, it is also useful to be able to examine a single, globally in-focus image of a non-planar scene.

After a stack is captured by the MI and a dense correspondence is found, images are transformed using the dense correspondence to create a new stack of images, where each pixel x,y on one image corresponds directly to the same pixel x,y on all other images. Next, these transformed images are combined into a single, in-focus image, called the focal section merge.

A simple approach to combining the images would be to compute each pixel in the focal section merge by selecting the pixel with sharpest focus from that location in the stack of images transformed to remove motion. Unfortunately, the transformed images have undergone interpolation from the original images, destroying some of the focus information. Our approach is, for each pixel in the transformed image, to use the inverse of the dense correspondence map, and use the corresponding location in the original image to compute sharpness..

Sharpness of focus $s(x,y)$, is calculated from local variation in image intensity $i_{x,y}$:

$$s(x,y) = \sum_{\substack{x'=x-3\sigma^2 \\ y'=y-3\sigma^2}}^{\substack{x'=x+3\sigma^2 \\ y'=y+3\sigma^2}} e^{\frac{(x-x')^2+(y-y')^2}{2\sigma^2}} \left((i_{x',y'} - i_{x'-1,y'})^2 + (i_{x',y'} - i_{x',y'-1})^2\right) \quad (10)$$

$$\sigma^2 = 3$$

After these sharpness values are computed, the focal section merge is constructed using the pixel from each location with the highest sharpness value. Figures X.a and X.b show the results of this technique for a stack (two stacks) of images from MER. Figure 8 shows a stack of two images, followed by the focal section merge.

3. 3-D Image Processing

Building 3-D models using stereo

Often the IDD arm was moved laterally to the camera pointing direction, offering views of a rock surface with sufficient parallax for recovering depth from stereo. The MI camera was very well calibrated before launch, but this calibration only applies to intrinsic camera parameters and the relationship between the camera coordinate frame and the IDD arm. Unlike the calibrated stereo camera pairs on the rovers, MI image pairs do not have exact calibrated extrinsic relationships. The IDD arm estimates the camera location to within about 1 mm, and this was often enough to recover epipolar geometry, then use the Ames Stereo Pipeline[citation] to perform dense subpixel stereo matching and 3D model creation. Figure 9 and Figure 10 show two MI images, and **Figure 11** shows two views of the mesh constructed from stereo.

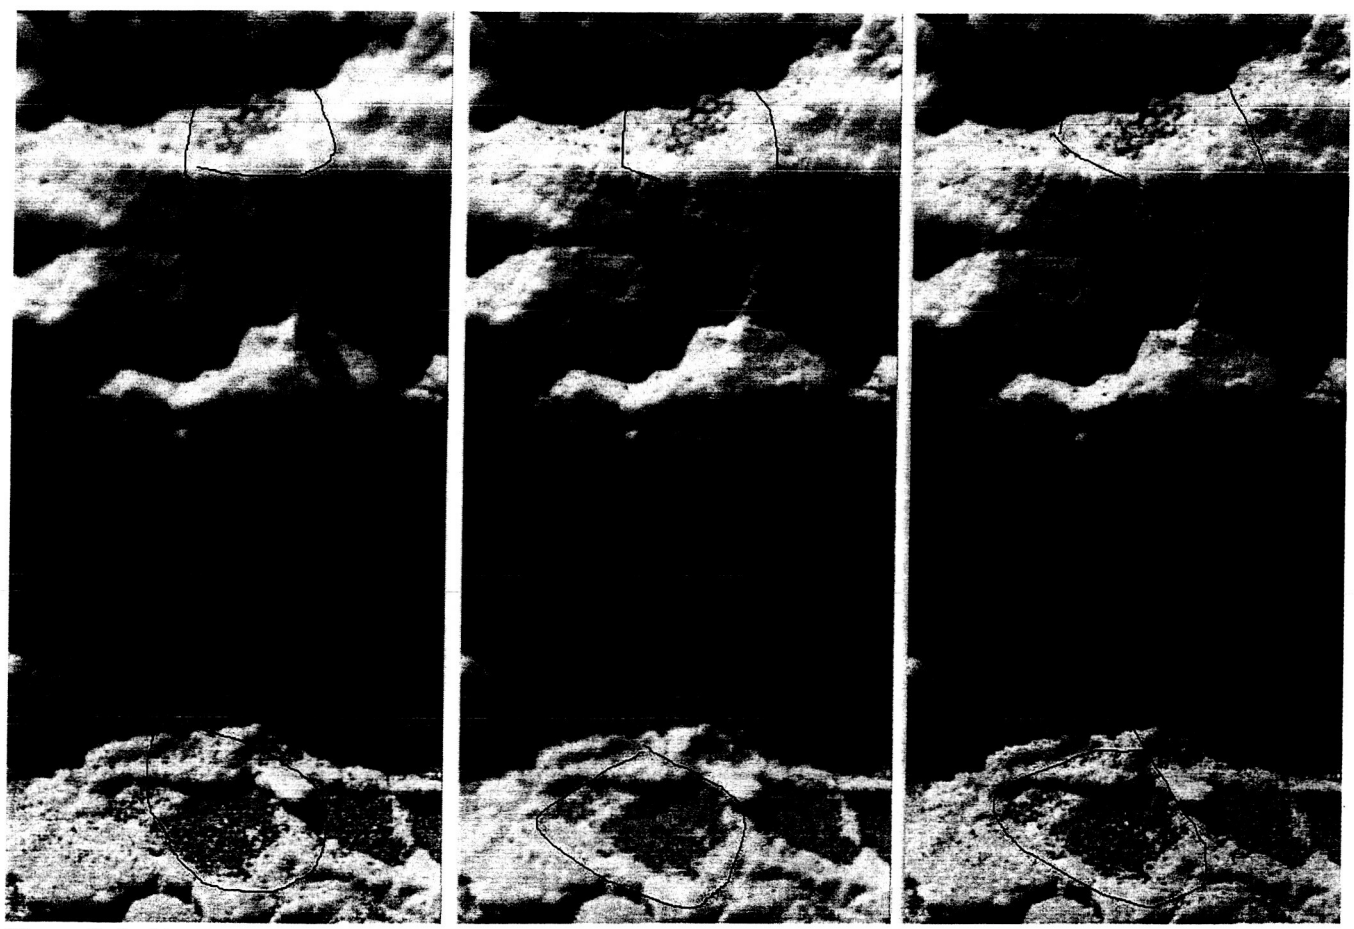

Figure 8 Left and middle images are original images from stack. Note that the upper circled area is out-of-focus on the left image, and the lower is out of focus on the right image. The right image is the focal-section merge, in which both areas are in-focus.

Figure 9 Left image of MI stereo pair of surface grinded by RAT. Notice the "blueberry" in the lower right quadrant of the image

Figure 10 Right stereo image. The berry is at the lower left edge of the image.

As the IDD moves to take the second image of a stereo pair, part of the first image moves out of the field of view. The resulting 3-D models include only the overlap between the first and second image. Typically during the mission image baseline was half the image size with no "toe-in", producing a model roughly 15.5 mm by 31 mm.

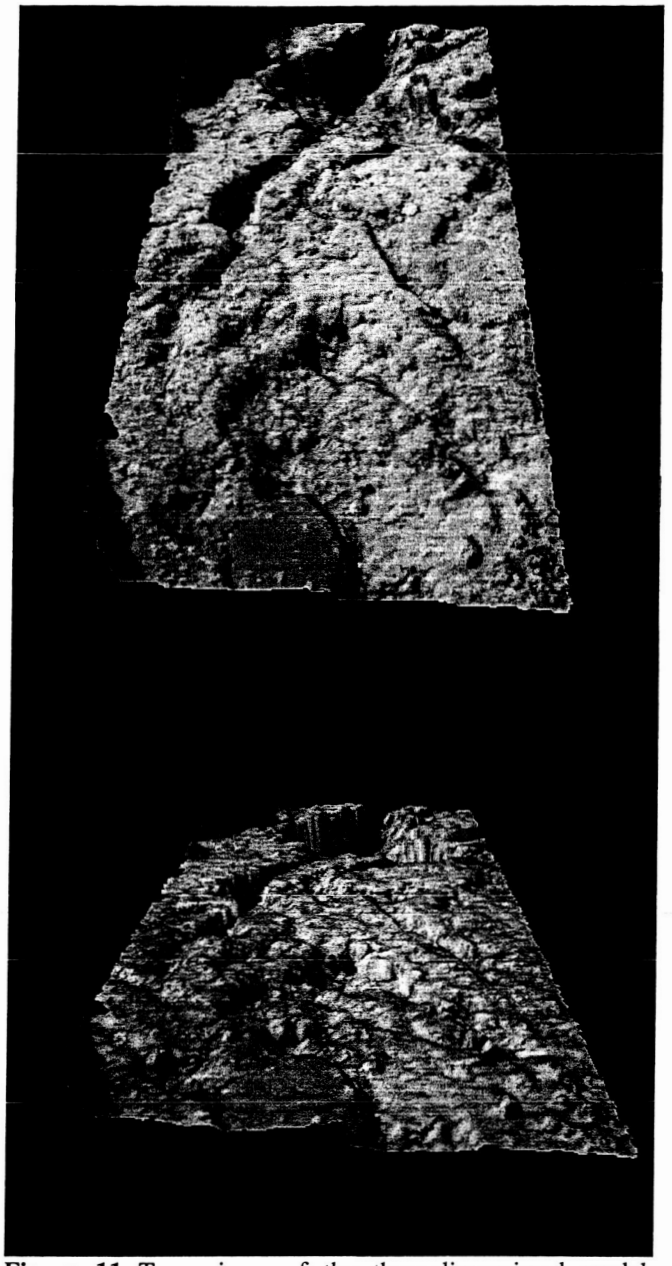

Figure 11 Two views of the three-dimensional model. Notice that the sliced berry in the foreground protrudes out from the surface. This protrusion is believed to have caused a RAT motor stall.

Building 3-D mosaics

Once a 3-D model is built, it can be viewed and manipulated using a tool such as Viz [9]. However, sometimes the MI stereo models are too small to cover an area of interest. Stereo pairs can be taken in a mosaic pattern, so that a high resolution 3-D reconstruction of a reasonably large surface can be built.

While the motion of the IDD is well known enough to provide constraints on the correspondence search for stereo, it is typically not well-known enough to automatically register multiple 3-D MI models to each other. The IDD's

motion is controllable to a minimum of 2±1mm RMS along the normal vector to the surface being imaged. An error of 1mm in the expected location of the camera could correspond to an error of 4.2mm in depth, assuming a 15mm baseline and approximately 63mm to the object being imaged. The error would be systematic across a single image pair, and the effect would not be noticed, but several different errors in depth would cause multiple 3-D models to not line up properly. The MI Toolkit incorporates mesh registration techniques developed at Ames [3] to compensate for these errors, though these techniques apply less effectively when the overlap between 3-D models is not large. Since the downlink bandwidth is limited, the overlap between successive 3-D models is usually insufficient for a fully automatic registration, and hand tuning is required. Figure 12 shows several MI meshes generated from a stereo mosaic of the surface of "last chance" registered to each other, and Figure 13 shows a similar stereo mosaic of the "dells" region.

Because the internal camera parameters are so well known for all of the imagers on board, 3-D registration is also possible between MI models and models taken with other camera pairs, particularly the pancams. Again, because of possible errors in IDD position estimation, translation and scale typically have to be hand-tuned to match the pancam or navcam mesh. Figure 14 is an example of an MI mesh of a RAT target embedded into a 3-D pancam model.

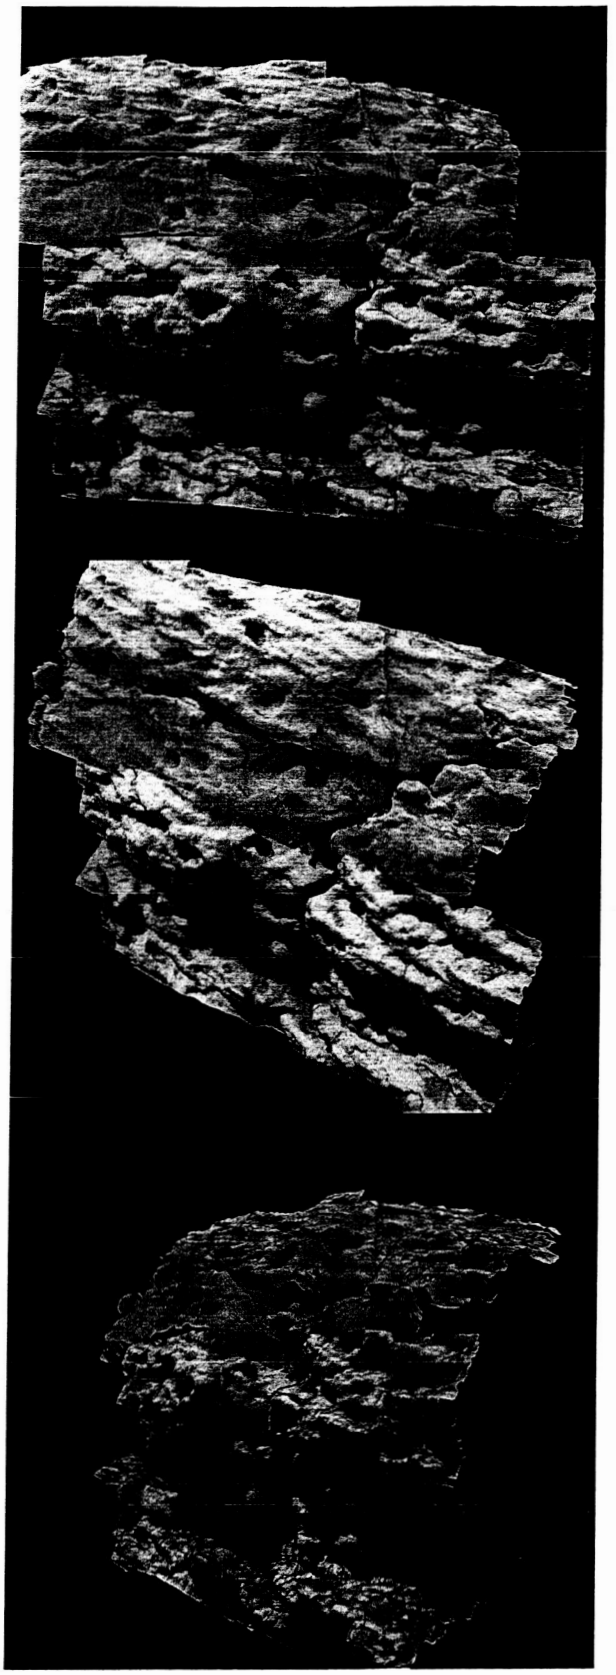

Figure 12 Three views of a MI stereo mosaic of the "last chance" rock. This is a view of ten individual stereo meshes, registered into a single model.

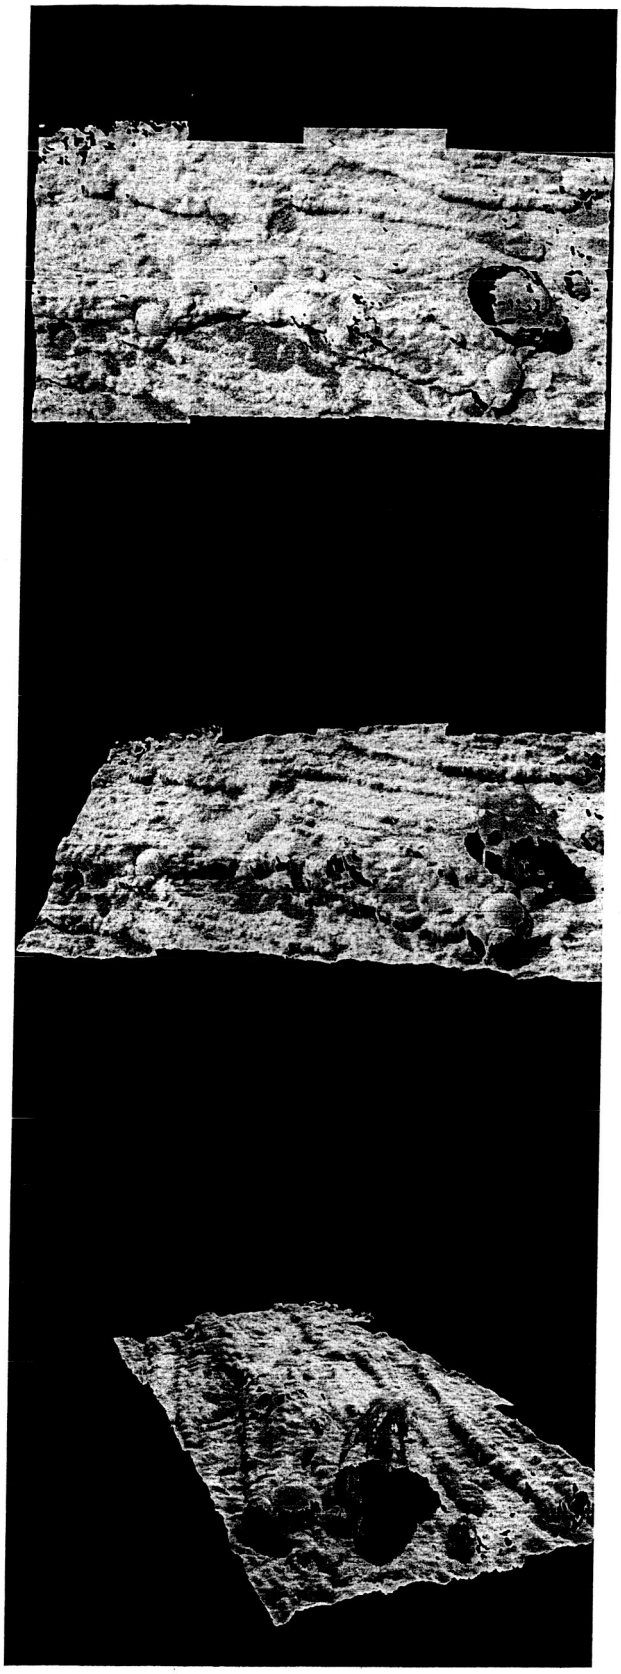

Figure 13 Three views of a MI stereo mosaic of the "dells" region. This composite model is a registration of four MI stereo meshes.

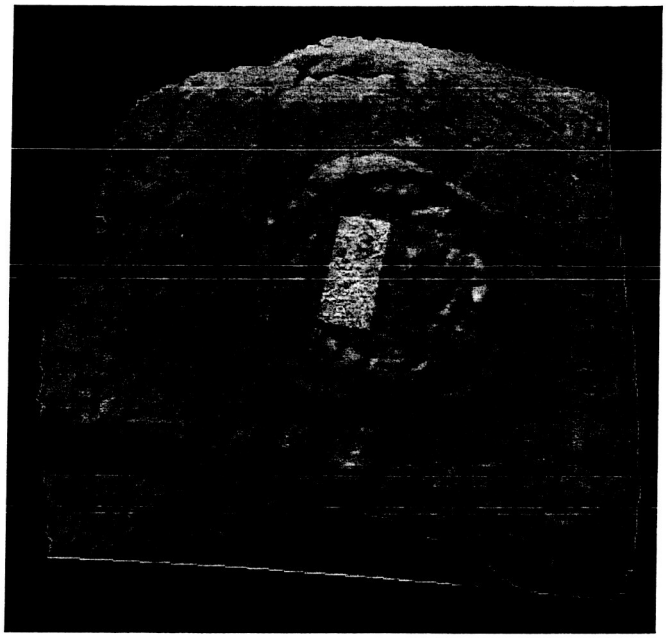

Figure 14 . The red region is a false-color pancam mesh; the gray region is the MI mesh

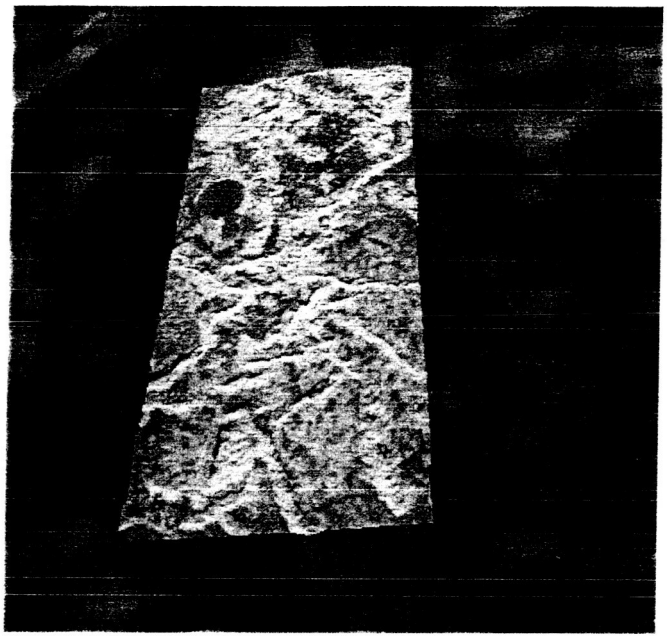

Figure 15 Close up of the previous figure. Note the difference in resolutions of the pancam and the MI.

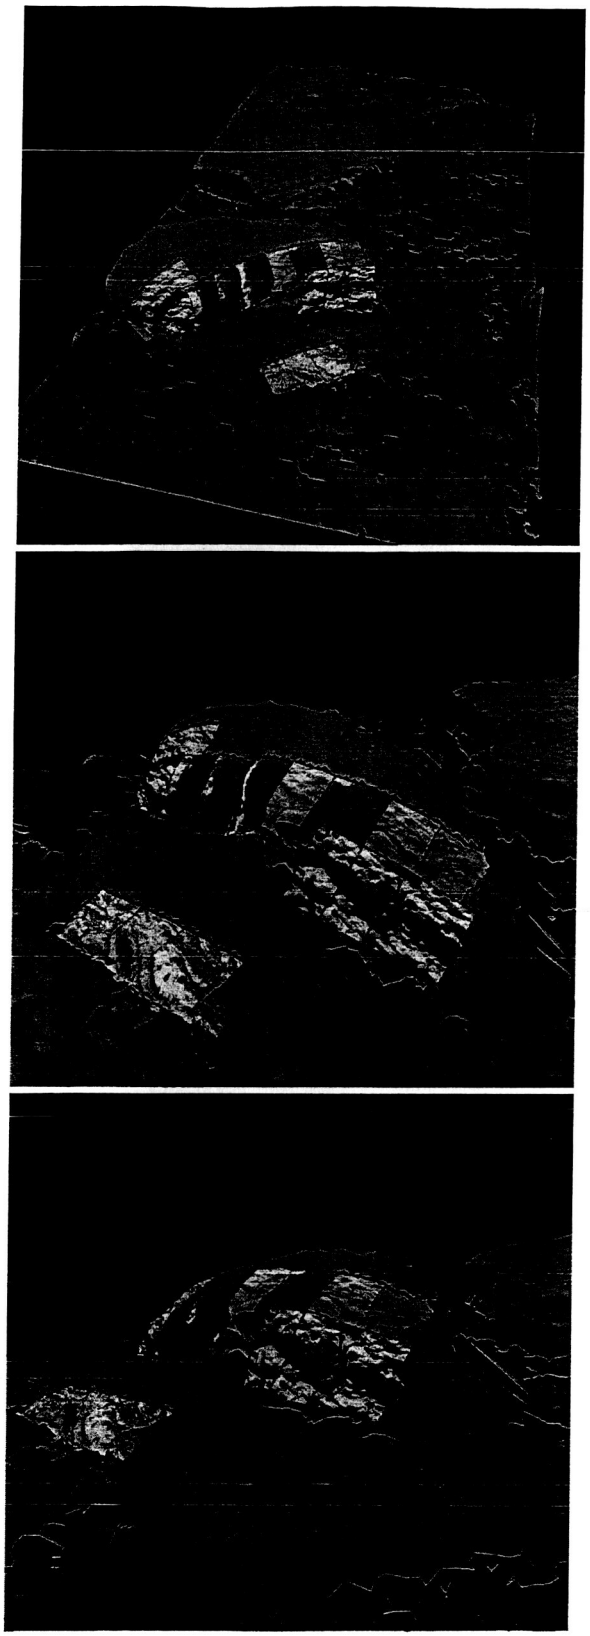

Figure 16 Three views of the "last chance" MI mosaic. The 3-D structure of this model, as well as the low-resolution texture map, are generated from pancam imagery; the high-resolution MI images are "draped" into the model.

Embedding 2-D into 3-D

When stereo pairs are not available for reconstructing 3-D models directly from MI images, the MI Toolkit can be used to embed MI images into 3-D models created using other cameras. While this does not provide scientists with the highest resolution possible in 3-D, it does provide a high resolution 2-D texture map for a lower-resolution 3-D model. This is particularly helpful for understanding the broader context in which the MI images were taken.

The locations of the pancams and navcams are well-known relative to the rover, allowing us to create a 3-D mesh using a stereo pair that is referenced to the rover's coordinate frame. Similarly, using the IDD kinematic model, the locations of the MI are approximately known at the time the MI images to be embedded are taken. For each MI image to be embedded, we copy the pancam or navcam mesh. Then, for each point in the original mesh, we use the MI camera model for each of the MI locations, projecting the point into the image plane of the MI. The first location that "sees" the point (the point projects to a pixel location that lies between 0 and 1024 on each axis) is then assigned that point, and it is removed from the other meshes. Finally, a texture map is built for each mesh using the pixels from the images that are seen by each camera. **Figure 16** shows one result of this technique.

4. Microsopic Image Browser

The primary access of the science team to the MI toolkit data was via the Ames MI Browser. The Ames MI Browser is a web based image data base displayer which allows easy access to the image and image format desired. The user chooses the rover and then the sol and image stack of interest as shown in Figure 17.

Once the rover, sol and stack are chosen then the user is presented with four choices: the stack of original images, the stack of registered images, a stack of registered images with disparity removed and a single image which merges the best focused elements from all images in the stack. See Figure 18.

Any of these stacks of images are resented on a single web page with the image displayed being dynamically displayed by moving the cursor over any of the magnifying glasses on the right (see Figure 19.)

This allowed quick, easy access to the wealth of MI image data and was well suited to the needs of the science team.

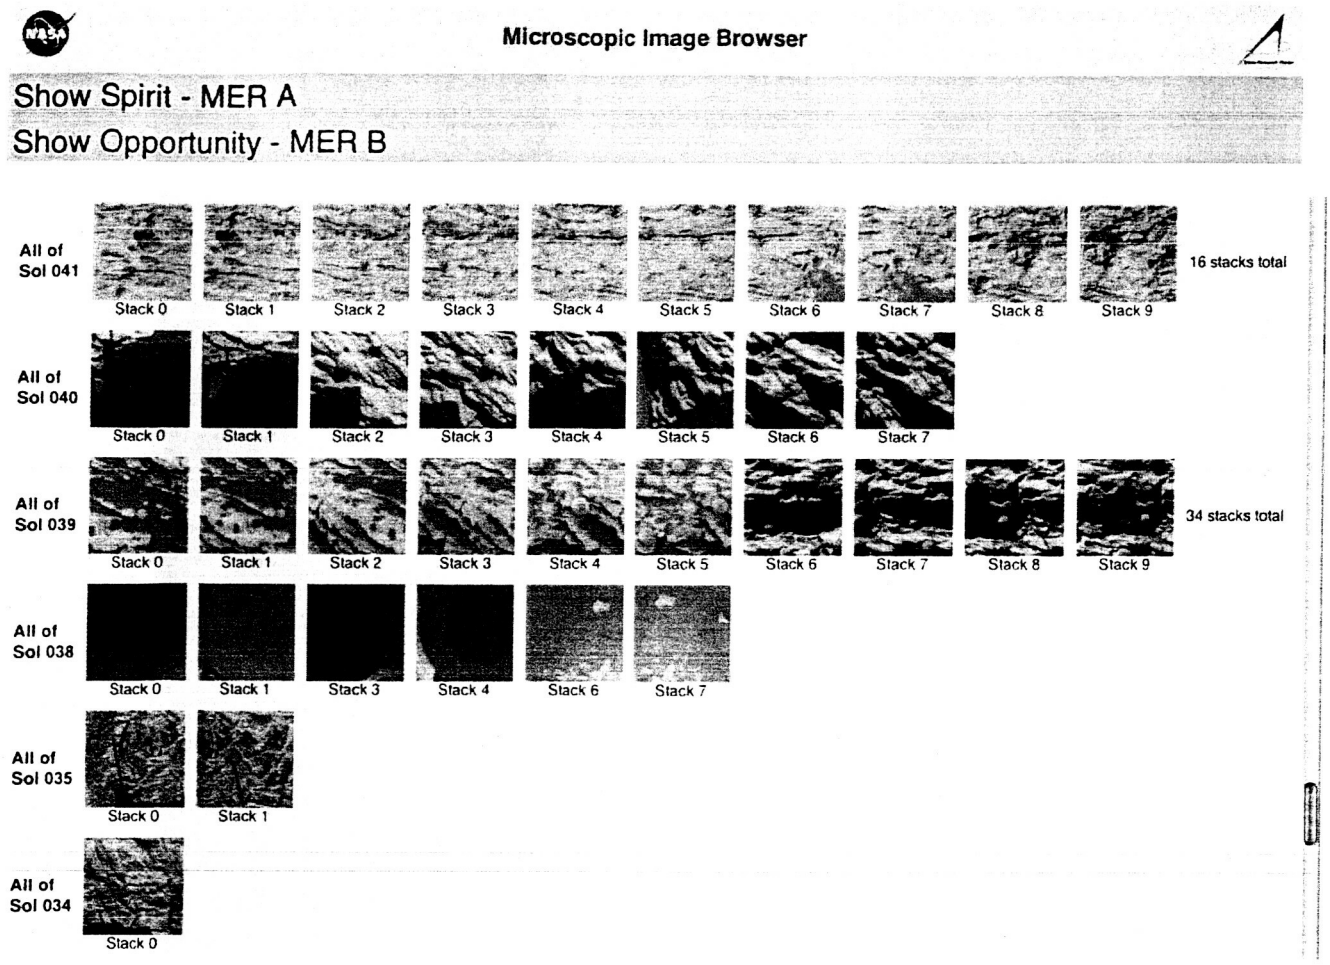

Figure 17 Initial interface for the Ames MIBrowser with choices of rover, sol and image stack

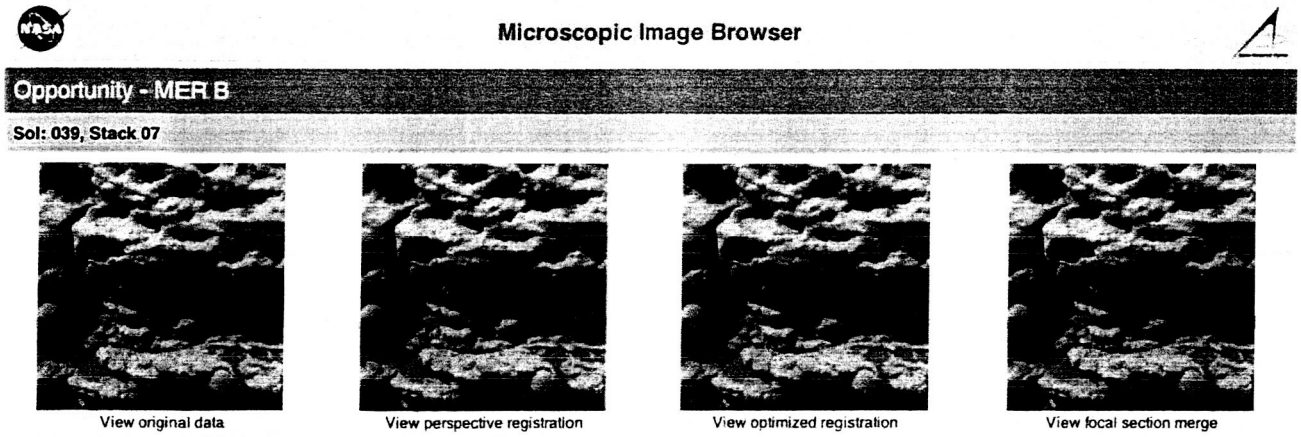

Figure 18 Stack of original images, the stack of registered images, a stack of registered images with disparity removed and a single image which merges the best focused elements from all images in the stack

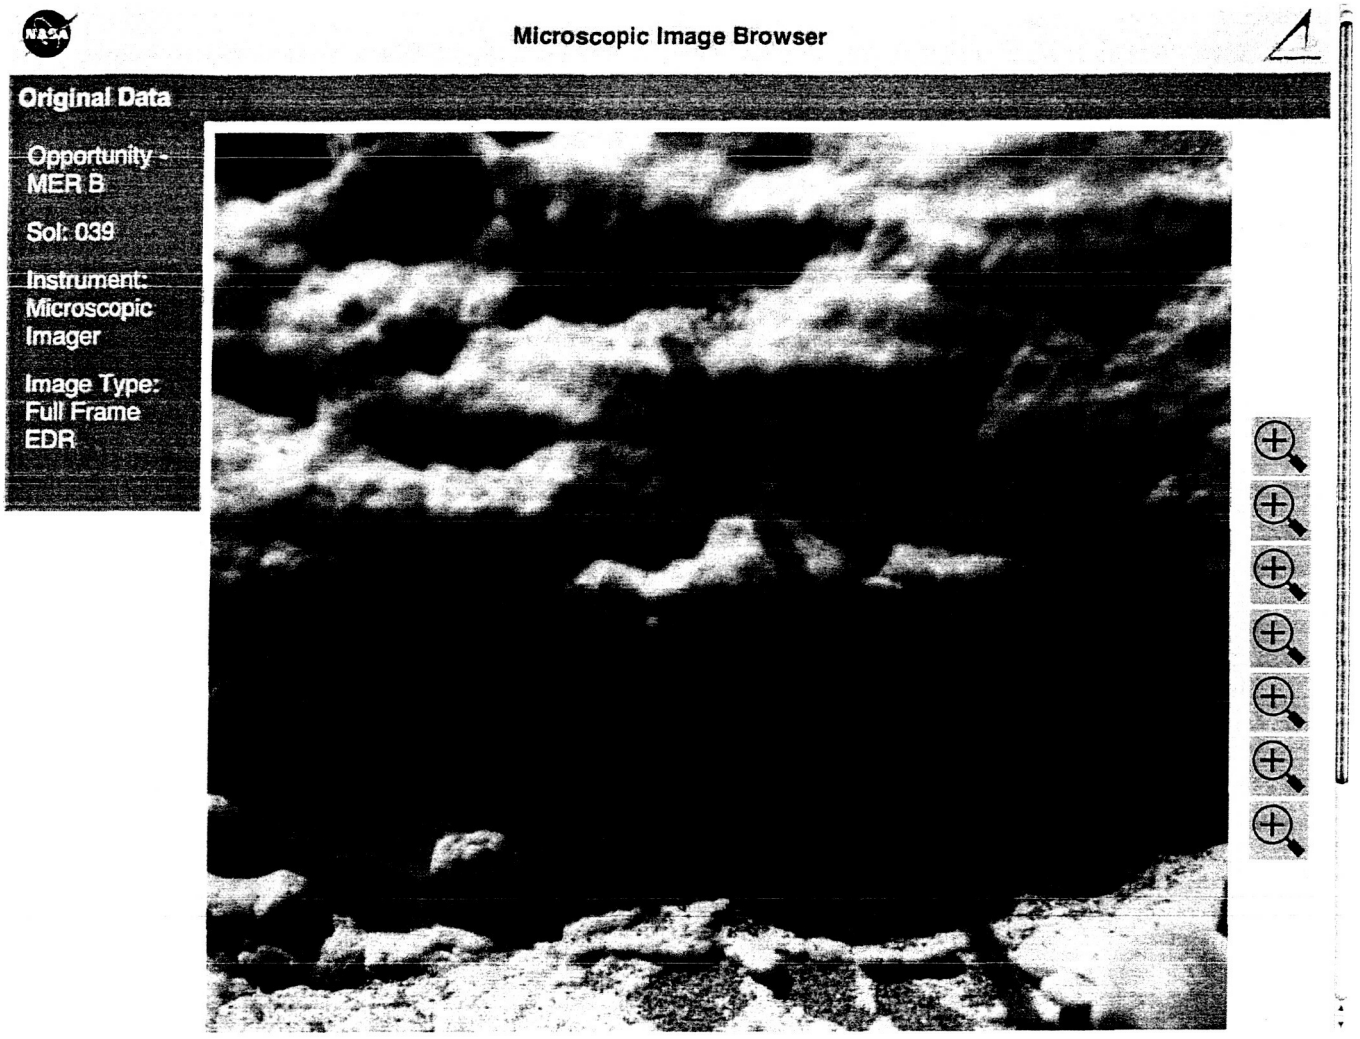

Figure 19 **User can quickly scroll through images in stack**

5. TOOLKIT USAGE DURING THE MER MISSION

The MER MI camera has repeatedly proved invaluable in the MER team's interpretation of Mars and the MI images were often little short of astonishing [8]. The textures and morphology at MI scales were unpredicted from the next best resolution images from the PanCam imager. From the perspective of the science team five issues are addressed by the Ames MI Toolkit:

- the rotation, scaling and alignment of stacks of MI images
- the 3D topography at MI scales
- mitigation of shadow effects, and
- the creation of a maximally focused single image from a stack of images with different in-focus and out-of-focus regions
- easy access to the raw and processed MI images.

After images are taken on a MER rover then they are downlinked to Earth either the following day or up to hundreds of days later. Images arrive one or more times during a day. Roughly once per day any new incoming MI images are sorted into stacks and within each stack they are registered, disparity is removed and the merging of focal regions. Then these images become part of the MI Browser's data base and then are available for evaluation by the science team. These images have been revisited over the entire mission and the tools of the Ames MI toolkit make this easier.

While we and our software were up to the tasks we expected from the mission, we found that we spent the majority of our time there trying to accommodate requests that were just outside the initial scope of the Toolkit, inevitably in situations where timeliness was extremely important. The versatility of the Toolkit and some rapid last-minute coding enabled us to provide timely data products for scientists to analyze.

6. ACKNOWLEDGEMENTS

R. Sargent, M. Deans, and C. Kunz gratefully thank the Athena science team for inviting us to participate in the exploration of Mars. We also wish to thank Larry Edwards and David Lees of the Viz team, and the rest of the K9 rover team at NASA Ames, including Maria Bualat, Mike Fair, Linda Kobayashi, Susan Lee, Eric Park, and Anne Wright. We finally wish to acknowledge Jimi Crawford, David Korsmeyer, Illah Nourbakhsh, Liam Pedersen, and the IS Program for their support of our involvement with this project.

REFERENCES

[1] J. Bilmes, A Gentle Tutorial on the EM Algorithm and its Application to Parameter Estimation for Gaussian Mixture and Hidden Markov Models. Technical Report, University of Berkeley, ICSI-TR-97-021, 1997.

[2] Y. Boykov, O. Veksler and R. Zabih, "Fast Approximate Energy Minimization via Graph Cuts," ICCV(1), pp377-384, 1999.

[3] Matthew Deans, Clay Kunz, Randy Sargent, and Liam Pedersen, *Terrain Model Registration for Single Cycle Instrument Placement*, Proceedings 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems, Volume 1, pp 323-328, Oct 2003.

[4] S. Geman and D. Geman. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images.*IEEE-PAMI*, 6, 1984, 721-741.

[5] D. Gennery. Calibration and Orientation of Cameras in Computer Vision, chapter Least-Squares Camera Calibration Including Lens Distortion and Automatic Editing of Calibration Points, pages 123-136. Springer Verlag (A. Gruen and T. Huang, ed.), 2001.

[6] D. M. Greig, B. T. Porteous, and A. H. Seheult. *Exact maximum a posteriori estimation for binary images.* Journal of the Royal Statistical Society, Series B, 51:271--279, 1989.

[7] K. E. Herkenhoff et al., Athena Microscopic Imager investigation, *J. Geophys. Res.*, 108(E12), 8065, doi: 10.1029/2003JE002076, 2003.

[8] K.E. Herkenhoff, et al., Textures of the Soils and Rocks at Gusev Crater from Spirit's Microscopic Imager. Science, Vol. 305:824--826, 2004.

[9] L. Nguyen, M. Bualat, L. Edwards, L. Flueckiger, C. Neveu, K. Schwehr, M. Wagner, E. Zbinden, "Virtual reality interfaces for visualization and control of remote vehicles," Autonomous Robots 11(1), 2001.

[10] C. Stoker, E. Zbinden, T. Blackmon, B. Kanefsky, J. Hagen, C. Neveu, D. Rasmussen, K. Schwehr and M. Sims, "Analyzing Pathfinder Data Using Virtual Reality and Superresolved Imaging", Journal of Geophysical Research - Planets Vol 104(E4) pages 8889-8906, April 25, 1999.

7. BIOGRAPHY

Randy Sargent *is software lead for the K9 rover target approach and instrument placement system at NASA Ames since 2002. In 1994 Randy co-founded Newton Research Labs, a machine vision company, with Anne Wright and Carl Witty. Randy led the Newton Labs team which won the 1996 and 1997 MIROSOT robot soccer tournaments, as well as the 1996 AAAI "Clean up the Tennis Court" robot contest. Randy left Newton Labs in 2000 to join Blastoff!, and in 2001 became Director of Open-Source Robotics at the KISS Institute for Practical Robotics. Randy holds BS and MS degrees from the Massachusetts Institute of Technology.*

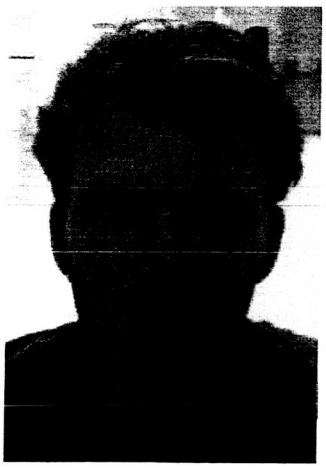

Matthew Deans *did his PhD work at the Carnegie Mellon University Robotics Institute, before joining the Intelligent Robotics Group at NASA Ames Research Center. He has developed simultaneous localization and mapping (SLAM) algorithms and sensor fusion based localization systems for rovers deployed in desert field sites in California, Nevada, Chile and Antarctica. He has also developed and supported machine vision ground tools used in MER mission science operations.*

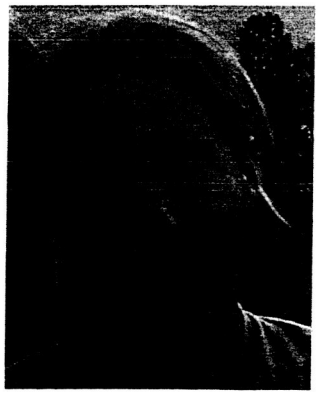

Clayton Kunz *is the lead software engineer for the K9 rover at NASA Ames, and spends much of his time working on computer vision problems for robotics. He is also the head of the math and data structures subgroup of CLARAty, a collaborative project developing a software architecture for robotic autonomy.*

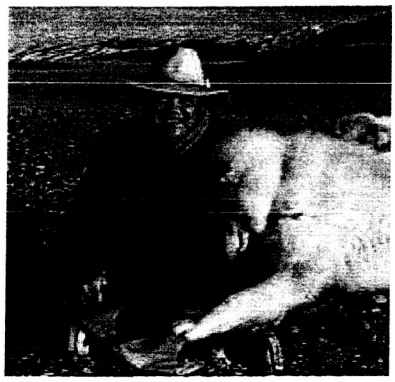

Michael Sims *received a BS in Physics and a Ph.D. in Computer Science and Mathematics from Rutgers University and has been at NASA Ames Research Center since 1987 and is currently a member of Ames' Computational Sciences Division and Center for Mars Exploration. His research includes robotics, machine learning, and tools for enhancing and easing scientific modeling. He was a member of the artificial intelligence group at Ames since its beginnings and lead the Intelligent Robotics group. He has served as agent for artifical intelligence and robotics for NASA's Office of Exploration (Code Z) which had responsibility for future +human missions to the Moon and Mars. Michael is actively involved in plans for future planetary missions including robotic activities and human settlements on the Mars and the Moon. For Mars Pathfinder he was the technical lead for the Ames virtual environment modeling system, and is currently a Co-Investigator on the Mars Exploration Rover (MER) mission.*

Ken Herkenhoff *is the Microscopic Imager Payload Element Lead for the Mars Exploration Rover mission. Since 1998 he has held the position of Geologist at the U.S. Geological Survey in Flagstaff Arizona. Ken holds an A.B. in geology from U.C. Berkeley, and M.S. in geology and Ph.D. in Planetary Science from California Institute of Technology*