

Stanford Artificial Intelligence Laboratory
Memo AIM-254

December 19 74

Computer Science Department
Report No, STAN-G-74-47 2

STANFORD AUTOMATIC PHOTOGRAMMETRY RESEARCH

by

Lynn H. Quam
Marsha Jo Hannah

Research sponsored by

National Aeronautics and Space Administration

COMPUTER SCIENCE DEPARTMENT
Stanford University



STANFORD ARTIFICIAL INTELLIGENCE PROJECT
MEMO AIM-254

NOVEMBER 1974

COMPUTER SCIENCE DEPARTMENT REPORT
STAN-CS-74- 472

Stanford Automatic Photogrammetry Research

by

Lynn H Quam

and

Marsha Jo Hannah

ABSTRACT: This report documents the feasibility study done at Stanford University's Artificial Intelligence Laboratory on the problem of computer automated aerial/orbital photogrammetry. The techniques investigated were based on correlation matching of small areas in digitized pairs of stereo images taken from high altitude or planetary orbit, with the objective of deriving a 3-dimensional model for the surface of a planet.

This research was supported in part by **the National Aeronautics and Space Administration's** Langley Research Center under Contract No. **NAS1-9682** and in part by the Advanced Research Projects Agency of the Office of Defense under Contract No. **DAHC-15-73-0435**.

The views and conclusions in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of NASA, ARPA, or the U.S. Government.

Reproduced in the **USA**. Available from the National Technical Information Service, Springfield, Virginia 22151.

I. Statement of the Problem.

The automated aerial/orbital photogrammetry problem is a subproblem of the more general stereo computer vision problem: given a suitable pair of images taken of the surfaces of one or more objects, derive a 3-dimensional model of said objects. In our **case**, there is one object, a portion of the surface of **Mars**, which we wish to describe in terms of 'deviations in elevation from a specified ellipsoid--the astronomer's model of Mars. This elevation information can be presented in several visual forms: we will usually present it as elevation pictures or elevation contour maps.

II. Approach.

Our overall approach is the same in effect, but slightly different in implementation from that taken by traditional analog photogrammetry. While it is possible to use the traditional contour-following methods in an automated system, it is more efficient to organize the task in a slightly different fashion.

Our first step in automated photogrammetry is to attempt to match (put into geometric correspondence) as many points as possible in the two images. The best method that we know for matching points is based on maximizing the normalized cross-correlation between areas centered on the points in the two images. For efficiency, it is necessary that these areas be as small as possible, but they must be statistically valid, that is, the areas must have significant information content above the known noise levels of the data.

We use normalized correlation (as opposed to other available measures of match) because this measure is the most immune to the effects of linear errors in the photometric models for the two images. The most common such errors are vidicon spatial calibration errors (shading) and photometric function differences, due to our having different views of a surface for which there is not an accurate photometric model.

The result of the matching process is a set of **parallaxes** (differences in position between pairs of corresponding points in the two images), which are directly related to the elevations of the 3-dimensional points represented by the pairs. The second step is to translate these parallaxes into elevations. This takes into account the pointing angle data for the spacecraft at the times the pictures were taken, the geometric vidicon distortion, and any correction assumptions which need to be applied.

The result of this second step is an elevation picture--an array of numbers representing the relative elevation at surface points. These pictures are usually smoothed somewhat to fill in holes in the elevation data caused by low information in the data or errors in the matching process. In the final step, the elevation picture is contoured at some specific interval to form the desired contour map.

III. Restrictions on the Image Data.

Certain restrictions must be placed upon the data we can handle. Some of these are necessary to ensure that the images are comparable and that corresponding areas in the two images will "look similar". Other restrictions are

necessary to ensure that the variations detected on the surface are due to elevation.

The matching process requires that the two images appear to have been taken at very nearly the same range, with no appreciable spacecraft roll between the two views. So that the scale is consistent over the images, we require that the visible angular portion of the spherical surface be small. This rules out oblique shots which cause extreme geometric distortion between the two views. And, of course, for elevation work it is best that the images be taken at equal but opposite view angles with respect to surface vertical.

It is necessary that the images show no appreciable differences in sun angle or in time, since these variations in the surface features make it extremely difficult to match corresponding areas. Most important, we must have images with sufficient information and a good enough signal-to-noise ratio so that matching is possible. Images with high noise levels, featureless images, and **images** whose only features are linear in the direction of the camera baseline present problems which the current state of the art in matching cannot handle.

Ideally, we would like controlled raw images--images taken on the same orbit with moderate, equal but opposite view angles, no roll, etc. The lunar photographs taken by the Apollo astronauts and digitized for computer use are an example of such controlled raw images. It is hoped that the Viking landing site verification images will also be of this variety.

If controlled images are not available, we can work with rectified **images**--orthographic projections based on best estimates of spacecraft positions and orientations at the time the images were taken. The Mariner 3 B-frames are an example of such data. However, we prefer to work with controlled raw images, since **any** transformation of the images has the effect of smearing out high-frequency information.

IV. Limitations on Results.

Since we know that there will be photometric errors and noise in the system, it is not practical to attempt to place single pixels into correspondence. The best we can do is to match small areas. The optimal size of these areas will vary inversely with the general information content of the images, i.e. the correlation area must be larger if the overall information content of the images is low.

This limits how accurately we can spatially resolve changes in the parallax. Clearly if we are computing correlation over an $n \times n$ target area, and we shift the target area by $n/2$ pixels in either the vertical or horizontal direction, we still have a 50% overlap with the previous correlation window. Since the two areas are not independent, the parallax measurement for that second target area will not be independent of the parallax of the first target area. Thus, our spatial resolution will be inversely related to the necessary correlation window size. This, as well as the efficiency consideration, urges the use of windows that are as small as possible.

The accuracy limits in generating an elevation map are also determined by

how well we can estimate the parallax between points (areas) in the two images, Clearly, since area correlation accomplishes a gross registration of the pixels within the areas, parallax estimates done by this method will be within a half of a pixel of the correct value. Interpolation of the correlation values can give a closer estimate of the parallax.

There will also be some regions in the images which cannot be matched at all. This can be due to locally low information in the images, information which is directed solely along the baseline direction, or changes between the two images other than those caused by surface elevation.

V. Algorithms.

We implemented our contouring algorithms in three separate parts, represented by three separate programs. These are described in the following sections.

A. Establishing Matches.

Intuitively, two points--one in the first image and one in the second--match if they both represent the same place on the surface of the planet. Determining computationally whether or not two points match can be done in a number of ways. Our favorite definition is that two points match if the normalized cross-correlation between the $n \times n$ areas surrounding the two points is a local maximum [Quam, 19713 and is sufficiently high [Hannah, 19743.

Our basic algorithm for establishing large quantities of matches begins by finding a starting match. While there are several methods by which this can be done automatically [Hannah, 1974], for this application it is done by hand, that is, the operator picks a starting point in the first image, locates its match, and gives this point pair to the program,

Given a starting pair for which the correlation is a local maximum and sufficiently high, our technique is to "grow" a region of pairs which show a local maximum at the same integer parallax as the starting match. This is done by pushing the starting point onto a stack, then for each point on the stack, checking whether its eight neighbors have been evaluated before or **also** show a local maximum at the present parallax. Pairs which represent a local correlation maximum are pushed onto the growing stack and marked as having been evaluated; these points and their parallaxes are also recorded on a disk file for later processing. Pairs which are not a local maximum at the current parallax are marked as being mismatches (to prevent the grower from looping infinitely) and are pushed on a "mismatch" stack for later treatment.

When all contiguous points which show a maximum at a given parallax are exhausted, our algorithm takes one of the point-pairs from the mismatch stack and performs a local search to maximize the correlation. If that correlation is sufficiently high, then a constant-parallax region is "grown" around that pair as described above. This continues until the stack of mismatched pairs is exhausted.

Usually, when the stack is exhausted, the relevant part of the surface **has been covered**. If this is not the case, the operator picks a new starting

match and repeats the above process until the surface is sufficiently covered,

This algorithm can be performed for every point in the picture, or it can be performed at every k -th point, resulting in a grid of points. For most of our applications, the algorithm is used in grid mode, with $k=n/2$ (n is the **diameter** of the correlation window). Illustrations 1 and 3 (All illustrations appear together at the end of this report.) show the results of such gridded region growing. In both cases, the upper left image is the first picture overlaid by a dot at the position of each window center which we were able to match, while the upper right image is the second picture overlaid by a dot at the position of each matching window center.

In order to ensure that the windows over which the correlation is calculated are statistically valid, the part of the region grower which tests for a local maximum first checks to see that the local information, as expressed by the local **variance**, is sufficiently high over the window. More recently implemented, although not in time for use on the examples in this paper, was a check to see whether the information present is highly directional parallel to the current best estimate of the camera **baseline**. If the information content is not satisfactory; the point is marked as being unmatchable.

An earlier version of this program would, when insufficient variance was found, enlarge the window until it contained enough information or exceeded a preset threshold. Consider, however, what happens when this algorithm tries to cross a large crater whose rim is clearly delineated but whose floor has no information. On the rim, where there is enough information, the algorithm would be happy; when it runs out of information on the floor of the crater, it begins to enlarge the window. If the floor is truly informationless, the window will continue to enlarge until it includes a portion of the rim, which provides sufficient information. If an attempt to match is done on this window, the information at the rim will dominate the low information in the crater itself, and the match will occur at the parallax of the rim, not the parallax of the crater floor--a bad match. While this window enlarging seemed to be a good idea in theory, it did not work in practice, and is no longer being used.

In automatic mode, the program does not stop to ask the operator if each proposed re-match for a mismatch is a good one. Instead, it applies some simple tests to determine if the proposed match is plausible.

Given a model of the spacecraft positions and orientations when the two images were taken and a particular target point, the position of its matching point in the second image is constrained to a line which has roughly the direction of the camera **baseline**. [Hannah, 1974] Furthermore, since the 3-dimensional point lies within a few kilometers of the planet's ellipsoidal ideal surface, the matching image point is further constrained to a small segment of the above-mentioned line.

A bad match is indicated whenever a matching point falls a significant distance off of this line segment. Also, since one would expect the surface of the planet to be fairly continuous, a sudden, large change in parallax is suspect, even though it lies along the baseline segment. Finally, the correlation at a proposed match must be sufficiently high, as determined by applying a

variation of the autocorrelation threshold test described in Hannah [1974].

B. Elevation calculations.

The elevation calculation program operates on a set of **paral**laxes, the output of the **preceeding** algorithm. These are recorded on a **di sk f i l e** as the integer co-ordinates of the **target point**, the integer co-ordinates of the **matching point**, and the nine correlations (the target correlated with the **match point** and the target correlated with each of the eight neighbors of the **matching point**) which prove that this pair is a match.

These nine correlations are used to locate the real (non-integer) co-ordinates of the **matching point**. What we wish to do is to approximate the correlation surface, which is known only at integer points, in order to locate its maximum. This will give us the "true" non-integer **matching point**, hence the **paral**lax.

Various schemes for this approximation have been tried, including fitting paraboloids by least squares techniques. A crude but expedient method is to fit 4 parabolas to the correlation data--one to each of the 4 triples of data points which include the center point--horizontally, vertically, and on each of the two diagonals. If the horizontal-vertical pair of parabolas show a pseudo-maximum in the same vicinity as the diagonal pair, then the locations of these two pseudo-maxima are averaged to yield the real co-ordinates of the **matching point**. If these two pairs of parabolas do not give consistent pseudo-maxima, then the point-pair is rejected as having an irregular correlation surface which cannot be modeled.

This technique leaves something to be desired, for we find that many of the "holes" in our elevation data are due to rejected correlation surfaces. However, **as** yet we have not come up with a better method which is as computationally expedient.

Once the real parallaxes have been determined, the task of converting them to elevations begins. If the positions and orientations of the spacecraft are accurately known, this is very simple. If, however, the spacecraft data is unknown or unreliable, as is the case with some of the Mariner 3 images, then a relative camera model must be derived from the parallaxes.

The first step in deriving a simple camera model is to approximate the **paral**laxes linearly, that is, to explain the **matching points (u,v)** from the **target points (x,y)** as

$$u \approx a*x + b*y + c \quad \text{and}$$

$$v \approx d*x + e*y + f$$

This is done by least squares techniques to minimize the residuals

$$(r,s) = (a*x + b*y + c - u, d*x + e*y + f - v)$$

over all of the **paral**laxes available. Next the baseline direction is fit to the

residuals (r, s) by finding the angle α which minimizes (least squares again) the quantity $(-r*\sin(\alpha)+s*\cos(\alpha))^2$, the off baseline distance of the residuals,

With this information, we start all over again. This time we do a simultaneous fit to the equations

$$u' \approx \cos(\beta)*x' + \sin(\beta)*y' + g \quad \text{and}$$

$$v' \approx -\sin(\beta)*x' + \cos(\beta)*y' + h$$

where (x', y') and (u', v') are (x, y) and (u, v) rotated through the baseline angle α --analytically putting the camera baseline along the x' axis. This least-squares fit calibrates the relative translation (g, h) and roll β between the two images. These last two steps (fitting the baseline and fitting the "camera model") can be iterated, if desired, to increase the accuracy of the model,

The final residuals--what is left after taking the translation, roll, and baseline angle into consideration--amount to a distance along the baseline and a distance off of the baseline. The distance off of the baseline is an indication of the accuracy with which the parallax can be determined. If, for any of the parallaxes, this component is too large, that parallax is rejected as being inaccurate.

The distance along the baseline is the elevation parallax. When multiplied by the appropriate conversion factors, so that it is expressed in meters of elevation on the surface, rather than pixels of parallax, this gives the relative elevation of that point in the image.

This elevation forming program receives parallax data in the scrambled order that the region growing program produced it. Under one option, it simply puts the data onto another disk file in the same order, as the elevations are determined.

Under the other option, elevations are put into an "elevation picture", then scaled so that they use the entire range of the pixel values available and occupy only positive values. Illustrations 1 and 5 each have one of these elevation pictures as their lower right image. Such elevation data can then be smoothed to fill in any small holes left by the region grower or emptied by a correlation surface which cannot be modeled.

C. Contouring.

We have used two contouring algorithms in our work. The first algorithm takes as input a rectangular array or picture of elevation data over some grid spacing in the pictures. The elevation values are integers greater than zero and are surrounded by a border of zeroes. The data array can contain "holes"--places where no elevation data is available--which are symbolized by elevations of zero. The second algorithm takes as input a list of integer positions and real elevations of points, which it manipulates into a net of triangles. What the first algorithm would see as small holes in the data are here covered by triangles; larger holes are usually skirted.

Both contouring algorithms are quite simple. We are given a set of contours to draw, expressed by a starting contour and a contour interval. Beginning with the lowest contour level specified, we scan the elevation data structure for a cell (a grid square in the first algorithm, a triangle in the second) through which our contour passes. We trace this contour, recording where it goes, until it is traced in its entirety. We continue scanning until all contours at this level have been traced, then begin scanning again for the next contour level.

Each contour is traced by examining cells of elevation data. When a contour goes into one of these cells, the algorithm moves around the edge of the cell in the clockwise direction, looking for a way out. Finding one leads it into another cell; the position of the exit point is found by linearly interpolating the elevation data along that side of the cell. This is repeated until termination conditions are satisfied.

For the first algorithm, the termination condition is that the original square is entered again from the original direction. This is possible because this algorithm sees holes and edges as places of very low elevation and continues to draw contours by them. This ensures that contours are closed curves (even though those parts of the contour which border holes and edges are invisible to everyone but the program), so the contour follower will eventually get back to its starting place. Saddle points get special attention, so that they always appear as two separate contours. The second algorithm sees holes and edges as the end of the net structure, so it terminates when it gets back to the starting point or falls off of any edge of the triangle net.

The lower left image of Illustrations 1 and 5, the lower right image of Illustration 3, and the enlarged Illustrations 2, 4, and 6 all show contour maps produced by the first program overlaid on the first pictures of the indicated pairs. Extremely short contours have been discarded as being noise,

All holes in the elevation data for the first algorithm and those holes which are too big for the triangles to bridge in the second algorithm are treated in the medieval manner--we leave them blank and attach a mental label "**Here be Dragons**".

Errors in camera models, pointing angles, etc. usually require that some manual adjustment of the regional slope be made. Consequently, both algorithms have a provision for adding a term of the form $a*x+b*y+c$ to the elevation data to accomplish this adjustment. The lower left image of Illustration 3 and the upper left image of Illustration 5 show how this is done. The operator manually chooses three points which are believed to lie at the same elevation (shown in the illustrations by a dot and the elevation which the computer found at that point). The computer then fits the appropriate plane and adjusts the elevation data as it contours.

Although it is possible to apply smoothing functions to make our somewhat angular contours more "intuitive," we have not implemented algorithms for this purpose.

VI. Suggested Changes and Improvements.

Our algorithms have certain built-in limitations, caused by the manner in which we decided to do things and the computer system on which we are working.

First of all, we have limited ourselves to images which are about 200 x 200 pixels in area. With pictures of this size, it is possible to get both images on our video output device at the same time, which allows us to see what the program is doing, a very helpful thing in experimental programming. Also, pictures of this size can be kept entirely in core with our program code without resulting in an abominably large core load which our time-sharing system will penalize by running infrequently.

Since the original data is on the order of 1000 x 1000 pixels, our limitation to 200 x 200 pixels means that we must either work with portions of the original **data** or work with spatially reduced versions of the original data (or some combination thereof). We have tried our techniques on both spatially reduced images and windowed images,

However, for the Viking mission, contour maps will need to be made for entire images. Our techniques (particularly the region growing) can be used on small pieces of images with the results "glued together" at some stage. Alternatively, a scan-strip approach can be implemented which does something quite similar to the region growing, except on a strip by strip basis.

Such an algorithm need only keep two strips, one which is as wide as the picture and n pixels deep (n is the window height) out of the first picture, and one as wide as the picture and $k*n$ pixels deep out of the second picture. Here k is a constant such that all of the matches to the strip out of the first picture are expected to lie within the strip out of the second picture. Having $k > 1$ is required whenever there is a relative roll between the two images or when the camera baseline is in any direction except horizontal.

VII. Results.

We have worked intensively with two pairs of Mariner 3 images and two pairs of digitized images taken by Apollo astronauts. Our overall success on this data has been less than spectacular.

Only one pair of Mariner 3 B-frames, 160814 (DAS=7326758) and 238803 (DAS=10132924), were anywhere near suitable for our task. Attempts to work with this pair, showing some of the Martian canyonlands, at a 2x2 spatial reduction were moderately successful, with some problems due to low information and linear edges along the baseline direction.

We sought to remedy the lack of information by working with parts of the image-pair at full resolution, a technique which has worked well on terrestrial images. The upper half of Illustration 1 shows this pair of images overlaid by the grid dots which indicate the matching obtained. The lower right image is the elevation picture derived from these matchings. The lower left image shows elevation contours at 400 meter intervals overlaid on the first image; this same data is shown enlarged in Illustration 2.

Careful examination of the contours will show that work on these higher-resolution images was not really successful. (For comparison, see the contour maps of this area done by Wu [1973].) There are several areas into which contour lines run, then stop. Most of these are areas of low information, and the amount of noise in the data was such that obtaining reliable correlations in these areas simply was not possible.

Several contours appear to be somewhat strange. In the area (165,135), for instance, contours display a funny hook. (Image co-ordinates begin at the upper left corner of the image and increase to the right for X, downward for Y. The small tick marks around the edges of the image indicate 5 pixels, medium-sized marks are 25 pixels, and large marks are at 100 pixel intervals.) This hook is not a real feature, but is due to the information content of the images being just slightly greater than the threshold, resulting in an unreliable match.

In the area (125,90) there are several contours which are very squished together. Again, this is not a terrain feature. Region growing proceeded toward this "cliff" from one side along a sharply shadowed ridge. The shadow edge gave the variance operator enough information to OK the area, but since this shadow lies in the direction of the camera baseline and has little information on either side of it, matching along it was not really valid. When a later region growing approached from the other direction, a very different parallax resulted, causing the apparent cliff. It was this example which prompted us to include the directionality operator in the list of tests on matching pixels,

Once low information areas and high contrast linear edges along the direction of the baseline are thrown out, the results look quite reasonable. Unfortunately, there are so few data points left that contouring becomes a guessing game. Matters were further complicated on this pair by changed sun angles, distortion resulting from extremely different view angles, as well as an abominable noise level. Therefore, we decided to waste no further time on this pair.

A pair of Mariner 3 A-frames, 109031 (DAS=5492373) and 146Y31 (DAS=6823913), were also attempted at full resolution. Like the previous **canyonlands pair**, these images of Nix Olympica were taken from somewhat different view **angles** and with different sun angles. When combined with the low resolution of the A-camera and the noise inherent in all of the Mariner 3 data, these factors made this pair difficult to work on.

The upper two images in Illustration 3 show the pair of pictures with dot overlays to indicate the matching obtained. On the lower left is the first image with the leveling points indicated; the lower right shows the contours which resulted. Due to the low resolution of the images, the contours are at 8 km intervals.

Oespite its overall size, Nix Olympica is a feature with very little local elevation relief. One would expect a contour map of it to show lots of concentric contour "rings" with very little evidence of broken terrain (see Wu [1973]). Thus one can see at a glance that most of the contours in the lower

right image of Illustration 3 are nonsense.

Running the region grower again with a higher minimum information threshold cuts out most of the jumbled matchings in the area to the left of the peak: however, this leaves a large hole in the contour map. As yet, the state of the art in computer matching in noisy images is not up to the ability of human matching in these same images.

We were fairly successful with the image pair with JPL numbers 185942 and 174659, film 0883, taken by the Apollo astronauts of a lunar peak. Our actual work was with 5 x 5 spatial reductions: the contoured image shown in Illustration 4 is twice this size, to make it easier to see the details.

Other than a few areas which were saturated in the digitization, most of the image matched up well. The contour map shown in Illustration 4 was done at intervals of one pixel in parallax: pointing angle data to relate this to elevations was not available at this writing.

For the most part, these contours appear to be quite reasonable. There appear to be a few minor glitches in contours. For instance, at about (300,110) there are some strange loops in the contour. We believe this to be an as yet unlocated bug in the contour drawing program, probably having to do with our handling of saddle points in the data.

We were less successful with the pair having frame numbers 2482 and 2481 taken by the Apollo 15 astronauts, and showing a lunar area of low relief. Because of the vast size of the digitized images (2000 x 2000), we chose to do a 2 x 2 spatial reduction to get the data into our computer from the tape. The pictures we worked with are shown in the upper half of Illustration 5. The lower right image is the elevation picture derived from the match data. The gradual gradient from light to dark in this elevation picture would indicate that the surface had a significant slope to it; however, we assumed from the look of the terrain that the area was flat. Therefore, we applied the leveling indicated in the upper left image by the overlaid points. When the data was contoured at 1/2 pixel intervals, the contour map of the lower left image was produced; the same data is shown in Illustration 6 in larger scale.

This area of the Moon has many small features which correlate well, resulting in a reliable, nearly complete mapping. However, the total difference in parallax for this pair amounted to 3.5 pixels. This means that any elevation contouring we would be doing would be based solely on information derived from interpolating parallax between pixels, something for which we have not found a completely satisfactory algorithm. Consequently, although many parts of the contours make some sense, the overall effect is chaotic, with contours cutting through craters and behaving in other strange fashions.

Summarized briefly, our results were as follows. On data of high information, high resolution, and low noise--such as the Apollo data--we were able to obtain reliable matchings for most of the area of the images. On less perfect data--such as the Mariner 9 data--reliable matches could be obtained for only parts of the images.

From the matchings, we could determine the elevation **parallax to within** half a pixel, but attempts to estimate the **parallax** more accurately were **not** always successful. From complete or nearly complete mappings, we could produce elevation pictures, "eye-ball" level the data, and generate contours. However, since the elevation data is only as good as the matching which produced it, the quality of the contour maps produced depends heavily on the quality of the images.

If the Viking images are nice, clean pictures like the Apollo imagery, then we can expect that a computer will do a fairly good job of producing contour maps for them, down to the resolution of a pixel in **parallax**. If, however, the Viking imagery is much like the Mariner 9 imagery in its noise characteristics, we do not hold much hope that a computer can produce highly accurate contour maps.




Illustration 1: These images are geometrical ly transformed, full resolution windows from the Mariner 3 B-frames 160B14 (DAS=7326758) and 238803 (DAS=10132924), showing some of the Martian canyonlands. The upper two pictures have overlays indicating the matching obtained. The lower right image is the elevation picture derived from these matchings. The lower left image shows elevation contours at 400 meter intervals overlaid on the first image.




Illustration 2. The same as the lower left picture of illustration 1, enlarged by a factor of 2.




Illustration 3. These images are geometrically transformed, full resolution windows from the Mariner 3 A-frames 109031 (DAS=5492373) and 146Y31 (DAS=6823913), showing Nix Olympica. The upper two pictures have dot overlays to indicate the matching obtained. On the lower left is the first image with the leveling points indicated: the lower right shows contours at 8 km intervals.




Illustration 4. The original images were JPL numbers 185342 and 174653, film 0889, taken by the Apollo astronauts of a lunar peak. This 5 x 5 spatial reduction of a window out of 185942 is overlaid with a contour map done at intervals of one pixel in parallax.




Illustration 5. The original images were frame numbers 2482 and 2481 taken by the Apollo 15 astronauts of a low-relief lunar area. The upper pair of images shows the 2×2 spatially reduced windows with which we worked: the left one is overlaid by the leveling points. The lower right image is the elevation picture derived from the match data. When this was contoured at $1/2$ pixel intervals, the contour map of the lower left image was produced.




Illustration 6. The same as the lower left picture of Illustration 5, enlarged by a factor of 2.

Bibliography

- Hannah, Marsha Jo [1974], "Computer Matching of Areas in Stereo Images", Ph.D. Thesis, Stanford University.
- Quam, Lynn H. [1971], "Computer Comparison of Pictures", Ph.D. Thesis, Stanford University.
- Wu, S. S. C., F. J. Shafer, G.M. Nakata, and Raymond Jordan [1973], "Photogrammetric Evaluation of *Mariner 9* Photography" in *Mariner Mars 1971 Project Final Report: Science Results*, Jet Propulsion Laboratory, Pasadena, California, pp 587-532.