*Computer Science*
*Technical Report*

Colorado State University

# Locating Shadows in Aerial Photographs Using Imprecise Elevation Data*

Mark R. Stevens  Larry D. Pyeatt  David J. Houlton
Michael E. Goss

E-mail: {stevensm|pyeatt|houlton|goss}@cs.colostate.edu

Technical Report CS-95-105

Computer Science Department
Colorado State University
Fort Collins, CO 80523-1873

Phone: (970) 491-5792  Fax: (970) 491-2466
WWW: http://www.cs.colostate.edu

# Locating Shadows in Aerial Photographs Using Imprecise Elevation Data*

Mark R. Stevens Larry D. Pyeatt David J. Houlton
Michael E. Goss

Computer Science Department
Colorado State University
Fort Collins, CO 80523-1873

E-mail: {stevensm|pyeatt|houlton|goss}@cs.colostate.edu

## Abstract

Realistic images of terrain can be created by texture mapping aerial photographs onto elevation data. However, lighting and shadows in the original photographs are a problem when multiple aerial photos are merged, or when the sun direction used for the generated image is different from the original aerial photograph. Our goal is to generate a terrain surface reflectance map from an aerial photograph. This will be used to simulate a new aerial photograph with different lighting and viewing directions. We are concentrating on forest and mountain areas for use in ecosystem management and modeling applications.

One key to generating realistic new images is finding shadows cast by the terrain in the original aerial photographs. Digital elevation values are uncertain due to vertical imprecision in the measurement process and coarse sample spacing. This uncertainty can result in substantial differences between calculated shadow positions and actual shadow locations in photographs when conventional shadow location methods are used. We have developed a technique which takes into account this uncertainty to determine the probability that any given point on the terrain is in sunlight or shadow in the aerial photograph. This shadow probability is crucial to accurate calculation of the reflectance of terrain points corresponding to the pixels of the original aerial photograph.



## 1 Introduction

Aerial photographs provide a valuable tool for analysis in fields such as ecosystem modeling and land use planning. A three-dimensional model generated by using an aerial photograph as a texture map over a terrain elevation model provides highly realistic views. Interactive manipulation and exploration greatly enhances the utility of such images, allowing the user to generate images of the 3D terrain model from a variety of viewing and lighting directions.

Lighting will not be consistent when a mosaic is created from multiple aerial photographs, since the sun direction will change from photo to photo. Even when using a single photo, the desired sun direction for a generated image may be different from the original aerial photograph. To use an aerial photograph to generate a view of terrain at a time of day and/or time of year different than that at which the photograph was taken, it is necessary to change the lighting in the photograph for the desired time. We are approaching this problem with a two-step process: the first step will remove the effects of direct lighting by calculating the reflectance of each surface, and the second step will recalculate the light reflected from each surface based on the new lighting conditions.

Shadow detection in the aerial photograph is a crucial part of the removal of lighting effects. Aerial photographs are frequently taken at a time near local midday in order to minimize shadows, but the shadows are never completely absent, especially at higher latitudes.

Digital elevation values are imprecise due to limitations of the measurement techniques and the coarse resolution of the samples. The vertical measurement techniques used have a certain error tolerance, and the horizontal sample spacing limits the size of the smallest features which can be captured. Any attempt to predict the location of shadows on the terrain must account for this uncertainty in the

terrain geometry. A small variation in an elevation can cause a large variation in the location of the edge of the shadow which is cast. Any inaccuracy in the predicted shadow location is further magnified by the discrepancy in resolution between the elevation data and the photograph. Available elevation data typically has a sample spacing of about 30 meters, while low altitude aerial photographs can commonly resolve detail to less than 1 meter.

Prior research in the area of shadow detection in aerial photographs has not considered the imprecision of the terrain elevations. Franklin [Fra91] locates shadows based on height fields using LANDSAT and SPOT satellite images, which have much lower resolution than low-altitude aerial photographs. Thirion [Thi92] uses a detailed 3D scene model to find shadows in scenes, employing a simple histogram technique to remove errors. The technique described is designed for scenes containing buildings and other regularly shaped and easily modeled artifacts, and is not directly applicable to use with imprecise terrain models. Other work such as [Zho92] ignores shadows by assuming a high sun angle. Techniques developed by Shu [Shu90] remove shadows of clouds in aerial photographs. Although we are not addressing cloud shadow removal, the detection of fuzzy and irregular edges should prove applicable to removal of terrain shadows with uncertain edge locations.

In this paper we describe a new technique for predicting the probability that any location on the terrain is in shadow for a given sun direction. We will use this *shadow probability map* to guide a search for shadow regions in an aerial photograph. Unlike other techniques, the shadow probability map can use low resolution terrain data to predict the location of shadows in an aerial photograph of much higher resolution.

# 2 Shadow Probability Map Generation

## 2.1 Terrain

The terrain is described by a height field (a rectangular grid of elevations). A common form is the U.S. Geological Survey (USGS) Digital Elevation Model (DEM) [Geo90]. In addition to the elevation samples, the DEM includes an RMS error value for the data set. In the mountainous areas with which we are primarily concerned, the RMS error value is typically on the order of 10 meters. This error value supplies the uncertainty in terrain elevation used by our algorithm .

## 2.2 Illumination

In order to determine shadow locations it is necessary to know the position of the light source. Based on input values specifying the date, time, and geographic location at which an aerial photograph was taken, we calculate the direction to the sun [PP76]. At present we model the sun as a directional light source; in the future we will extend our algorithm to allow for an area light source. We also ignore atmospheric refraction of sunlight, since this has a significant effect only when the sun is very near the horizon, and we do not expect to encounter this situation very often. In fact, most aerial photographs are taken near midday, when the sun is closest to zenith and atmospheric refraction is negligible.

## 2.3 Generating the Shadow Probability

In order to determine the probability that a point is in shadow, we cast a ray from each elevation point in the height field towards the sun. Our ray casting technique differs from traditional terrain rendering ([Mil86, K$^+$89, PP92], etc.) in that we generate a probability that an intersection occurs instead of a definite intersection.

Since the terrain height field is represented as a grid of rectangular cells, we can calculate the entry and exit points for every cell the ray intersects. The two intersection points are averaged together to obtain an intersection height for the cell. A probability function based on the height and the distance travelled through the cell is used to determine the probability that the cell does not cast a shadow on the ray's origin. A cumulative probability for all points of intersection is maintained, ultimately containing the probability that the original point is not in shadow.

### 2.3.1 Probability Function

Each point in the height field contains some amount of potential error. We assume that the error is normally distributed with a mean of zero and standard deviation of $\sigma$. To avoid excessive computation when a ray is far above or below the recorded elevation, we clip the distribution at one standard deviation above or below the sample point.

The USGS supplies an RMS error $\epsilon$ in the header of a DEM data set. We use this RMS error value as an approximation to the standard deviation $\sigma$ of our error distribution. For $n$ samples, we use the calculation

$$\sigma = \frac{n}{n-1}\epsilon \tag{1}$$

Note that $\sigma \approx \epsilon$ for large $n$.

When the ray passes through a cell in the elevation map, the height of the elevation map at that cell is subtracted from the height of the ray and the resulting distance metric is used to calculate the probability $p(c)$ that the ray does

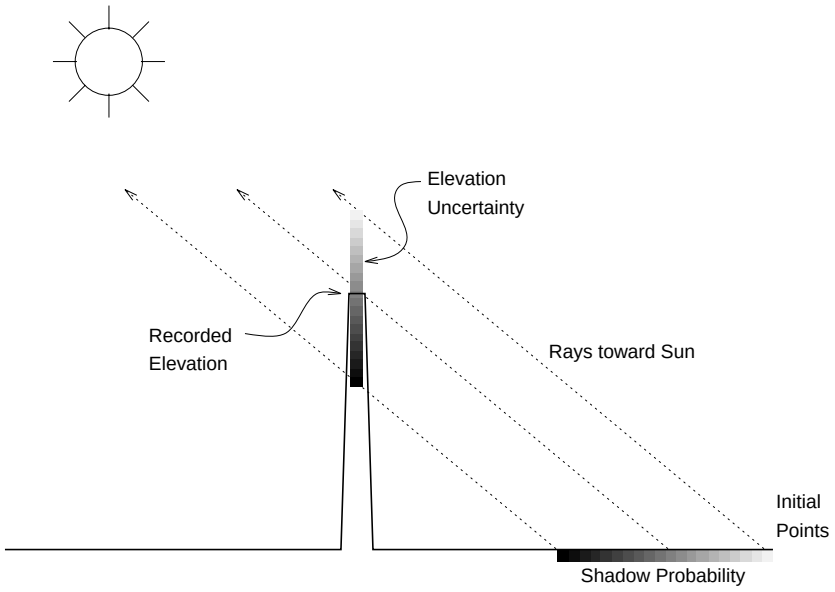


Figure 1: Effect of Elevation Uncertainty on Shadow Determination

not hit the terrain at that cell (see Figure 1). The probability is calculated using the cumulative distribution function for a normal random variable with variance $\sigma^2$ and mean 0:

$$p(c) = \int_{-\infty}^{z_{rc}-z_{ec}} \frac{1}{\sqrt{2\pi}\sigma} e^{-y^2/2\sigma^2} dy, \qquad (2)$$

$z_{rc}$ is the height of the ray at cell $c$ and $z_{ec}$ is the height of the elevation map at cell $c$. The value obtained from equation 2 represents the probability that the current cell does not cast a shadow on the cell from which the ray originated. If we let $C = \{c_1, c_2, c_3, \ldots, c_n\}$ be the set of all elevation cells through which the ray passes, and $D = \{d_1, d_2, d_3, \ldots, d_n\}$ be the distance that each ray travels inside each cell relative to cell size, then the cumulative probability is given by

$$P = \prod_{i=1}^{n} p(c_i)^{d_i} \qquad (3)$$

$P$, the shadow probability map entry, represents the probability that the original cell is not in shadow from any cell.

### 2.4 The Shadow Probability Map

A grey-scale image can be generated to illustrate the shadow map (see Plate 2 for an example). Lighter pixels correspond to shadow probability map values close to 1.0, representing areas unlikely to be shadowed, while dark pixels represent areas likely to be in shadow (probabilities close to 0.0).

## 3 Results

The results of the shadow probability map algorithm are illustrated using two different methods:

1. A shadow probability map generated from elevation data is compared to an aerial photograph.

2. The effect of uncertainty on the shadow calculation is illustrated by subsampling elevation data and adding noise.

### 3.1 Aerial Photograph

To demonstrate the algorithm described above, a shadow probability map was generated for the Clark Peak area of the Colorado State Forest. The USGS DEM grid for this area corresponds closely, but not exactly, to the aerial photograph shown in Plate 1. This is a section of the original photograph; it has not been warped to create an orthophoto. The photograph was taken on September 4, 1988, at 10:54 a.m. MST. The elevation data consists of 7.5 minute USGS DEM sampled on a 30 meter grid, centered at UTM 423300 easting and 4497390 northing in zone 13. The elevation data is Class 1 with a measured RMS error value of 10.0 meters.

The photograph was taken near midday, and the high sun angle generates few shadows. There are, however, some areas on the north side of the steep ridge near the center of the photo that are shadowed by the ridge (there are several lakes near the top of the photograph that should not be mistaken for shadows). We first attempted to detect the shadows using a conventional ray casting technique, with

no uncertainty in the terrain elevation values. This technique generated no shadows anywhere in the image (a binary shadow map would have been completely white).

As can be seen in Plate 2, our shadow probability map technique indicates a strong likelihood (dark areas) of shadows in the areas which appear to have shadows in the photograph. The steep slopes near the bottom of the image have light gray values, indicating that minor variations in elevation from those in the DEM could also cause shadows in these areas.

## 3.2 Subsampled Data

The aerial photography currently available to us has limited numbers of shadowed areas due to the high sun angles. We expect the shadow probability map technique to also work with lower sun angles, and have generated some synthetic images to demonstrate this.

We start with existing 30 meter DEM data for the Richthofen USGS quad in the Colorado State Forest. For the purposes of the demonstration, this initial data is assumed to be error-free actual terrain. We simulate the terrain sampling process by subsampling this data to 1/4 the original dimension to create a new height field. At each point in the subsampled data set, normally distributed noise is added with a mean of zero and standard deviation of 10.0 meters (similar to the error in the Clark's Peak DEM). This process simulates point sampled terrain with a known "measurement" error.

Simulated aerial photographs of the original ("error-free") data were generated using a conventional ray tracing program (*Rayshade*). Plate 3 shows such an image, generated using the same date and time as that used for Plates 1 and 2.

Plates 4 and 6 show ray traced images of this same "error-free" data corresponding, respectively, to late afternoon (4:00 p.m.) and early morning (8:00 a.m.) on the same day. Next to each of these two images is shown the corresponding shadow probability map generated from the subsampled, noisy data (Plates 5 and 7, respectively). Notice that the hard shadows of the "error-free" terrain fall well within the high probability (dark) areas of the corresponding shadow probability map.

# 4 Future Work

## 4.1 Improvements

The results described above demonstrate that the shadow probability map is a useful tool for determining where shadows are likely to occur in aerial photographs. We foresee a number of improvements to the basic algorithm presented here, both in speed and in quality of results.

### 4.1.1 Quality

Some approximations currently limit the accuracy of the shadow probability map:

- The sun is treated as a directional light source; in actuality, it subtends a noticeable angle in the sky, and should be treated as an area light source.
- The uncertainty in the starting elevation of a ray is currently not taken into account.
- Since a single ray is cast from each point, some aliasing occurs in the sampling process.

As we integrate the shadow probability map into the shadow removal process, we will determine if the current accuracy is sufficient to guide the shadow removal process. Accuracy of the map can be increased by using multiple rays for each location. A stochastic sampling process [Coo86], jittering the ray direction and starting position, will account for the factors mentioned above.

### 4.1.2 Speed

The ray casting method described can be time-consuming for a large elevation map. An alternate method would use an orthographic projection of the terrain onto a viewplane perpendicular to the direction to the sun. This projection yields a "sun's-eye view" of the terrain, in which case surface shadow determination is equivalent to the well known computer graphics problem of hidden surface determination.

Following transformation of the data into viewing coordinates, a modified version of the floating horizon visibility test [Rog85] would be applied to each cell in turn. A shadow probability is assigned to each cell based upon its projected distance above or below the current horizon. Since each height field cell is considered only once, the computational cost of this method is low with respect to ray casting, and grows linearly with the size of the height field.

The shadow probability distribution produced using this projection method is similar but not identical to that produced by the ray casting method. Essentially, the ray-casting method ignores the error at the point of interest, generating a shadow probability by accumulating the error values of all the potentially occluding points along the path to the sun. The projection technique, on the other hand, considers the occluding points (the horizon) to be error-free, and calculates a shadow probability based solely on the error value of the point of interest. A more accurate approach, of course, would consider all possible sources of error in the probability calculation. A variation of the projection method which accomplishes this is currently under development.

### 4.2 Shadow and Illumination Removal

Determination of the shadow probability map is the first step in the process of extracting reflectance information from aerial photographs. Once the probability map is generated, edges of shadows must be accurately located. In HSV or similar color space, the shadow areas should appear as islands of different brightness (Value) within larger areas of similar Hue and Saturation. We will locate these islands using the shadow probability map in conjunction with a form of region segmentation, similar to that discussed in [BGK+89].

Once the shadow regions are located and labelled, we will proceed to solve for the reflectance value. Most terrain in the photographs used will roughly approximate Lambertian reflection. We will use USGS hydrography and land use/cover data to identify approximate locations of exceptions such as bodies of water. Classification based on the source images will refine the boundaries to the nearest pixel.

To find a terrain reflectance value corresponding to a pixel in the aerial photograph, we will use a standard Lambertian illumination calculation, solving for the reflectance term. The process above will give us appropriate input values for direct light (and occlusion), surface location, and resulting color (from the pixel of the photo). Ambient light from the sky will be added depending on the time of year and time of day. Surface orientation will be calculated using methods described in [GP93, Gos93]. The result of this process will be reflectance values corresponding to every pixel in the aerial photograph, suitable for use as a texture map with new illumination parameters.

## 5 Conclusions

Terrain rendering methods previously developed treat the underlying height field as error free. This approach is not suitable to our task of shadow removal. Therefore, we have developed an algorithm which correctly accounts for the errors inherent in a DEM derived height field which make exact shadow location difficult to calculate. This probabilistic approach is quite good at predicting areas where shadows are likely to occur. These predictions can be used to find and remove the actual shadows from aerial photographs. Comparison of our results with actual photographs and with ray traced images shows that the probability fields generated closely match the actual shadow areas.

This algorithm should prove valuable for achieving our overall goal of extracting reflectance information from a terrain photograph. Future work will concentrate on removing shadows and other illumination effects. The resulting reflectance values can be used to generate realistic images of the terrain from different viewpoints and under different lighting conditions.

## 6 Acknowledgements

Support for David Houlton's participation in this research is provided by a NASA Graduate Student Fellowship in Global Change Research. The aerial photograph in Plate 1 was provided by the Colorado State Forest Service and the Forest Sciences Department of Colorado State University.

## References

[BGK+89] J. Ross Beveridge, Joey Griffith, Ralj R. Kohler, Allen R. Hansen, and Edward M. Riseman. Segmenting images using localized histograms and region merging. *International Journal of Computer Vision*, 2(3), 1989.

[Coo86] Robert L. Cook. Stochastic sampling in computer graphics. *ACM Transactions on Graphics*, 5(1):51–72, January 1986.

[Fra91] Steven E. Franklin. Image transformations in mountainous terrain and the relationship to surface patterns. *Computers & Geosciences*, 17(8):1137–1149, 1991.

[Geo90] U.S. Geological Survey. *Digital Elevation Models*, volume 5 of *Data Users Guide*. Department of the Interior, Reston, VA, 1990. 2nd. (revised) printing.

[Gos93] Michael E. Goss. Enhanced shading of terrain surface data using object–space digital filtering. In *Proceedings of ETCE Computers in Engineering Symposium*, pages 163–172. ASME, 1993.

[GP93] Michael E. Goss and Ivor P. Page. Normal vector generation for sampled data using fourier filtering. *The Journal of Visualization and Computer Animation*, 4(1):33–49, January-March 1993.

[K+89] Kazufumi Kaneda et al. Three dimensional terrain modeling and display for environmental assessment. *Computer Graphics*, 23(3):207–214, July 1989. SIGGRAPH Conference Proceedings.

[Mil86] Gavin S. P. Miller. The definition and rendering of terrain maps. *Computer Graphics*, 20(4):39–48, August 1986. SIGGRAPH Conference Proceedings.

[PP76] G.W. Paltridge and C.M.R Platt, editors. *Radiative Processes in Meteorology and Climatology*. Elsevier Scientific Publishing Company, 1976.

[PP92] David W. Paglieroni and Sidney M. Petersen. Parametric height field ray tracing. In *Proceedings: Graphics Interface '92*, pages 192–200, May 1992.

[Rog85] David F. Rogers. *Procedural Elements for Computer Graphics*. McGraw–Hill, New York, 1985.

[Shu90] Joseph Shou-Pyng Shu. Cloud shadow removal from aerial photographs. *Pattern Recognition*, 23(6), 1990.

[Thi92] Jean-Philippe Thirion. Realistic 3D simulation of shapes and shadows for image processing. *CVGIP: Graphical Models and Image Processing*, 54(1):82–90, January 1992.

[Zho92] Qiming Zhou. Relief shading using digital elevation models. *Computers & Geosciences*, 18(8):1035–1045, 1992.

Plate 1: Clark's Peak aerial photograph

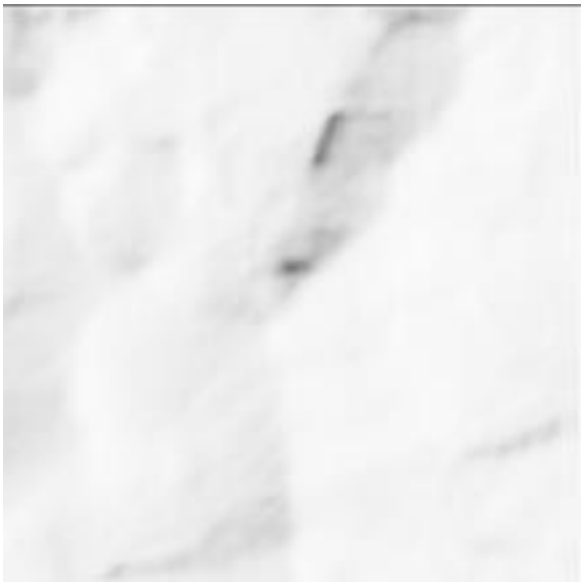

Plate 2: Shadow probability map for Clark's Peak

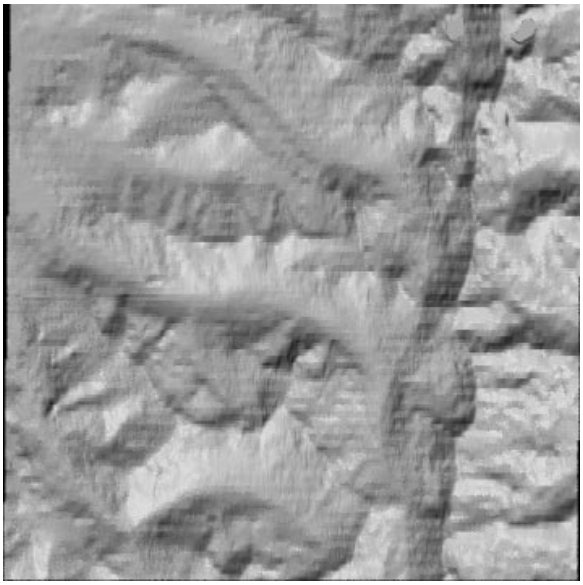

Plate 3: Ray traced image of Richthofen data (near midday)

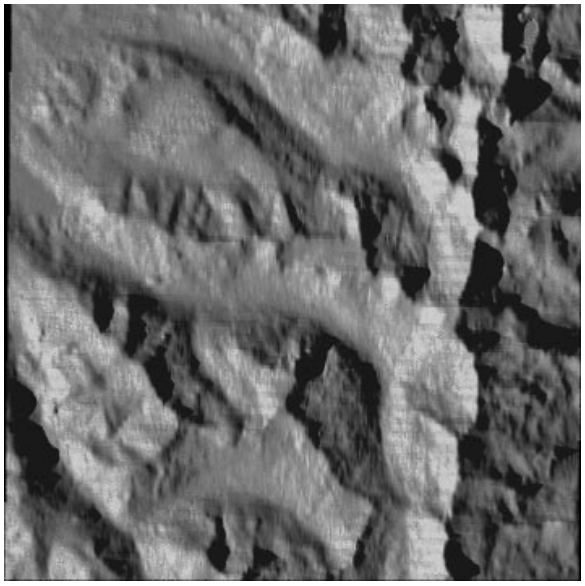

Plate 4: Ray traced image of Richthofen data (late afternoon)

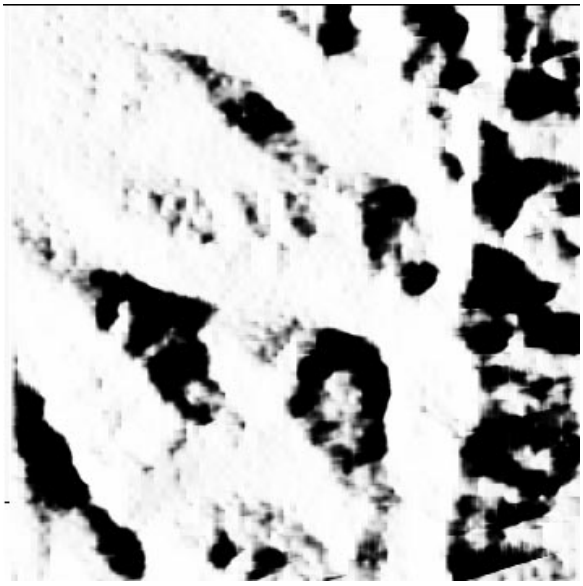

Plate 5: Shadow map for subsampled Richtofen data (late afternoon)

Plate 6: Ray traced image of Richthofen data (early morning)

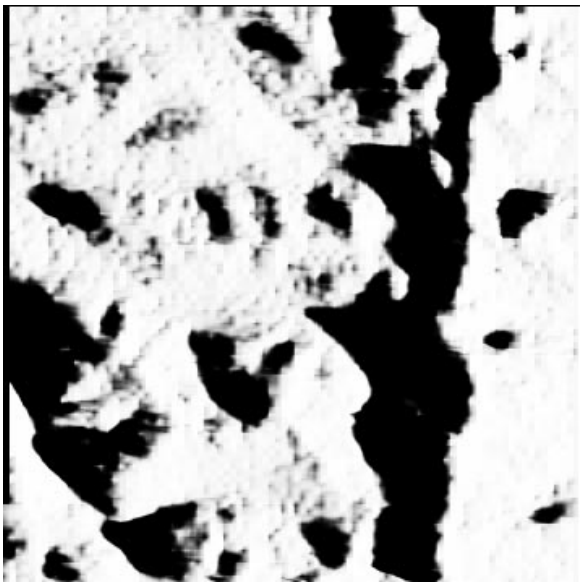

Plate 7: Shadow map for subsampled Richthofen data (early morning)