

AN OPTICAL MODEL FOR IMAGE ARTIFACTS PRODUCED BY DUST PARTICLES ON LENSES

Reg G. Willson, Mark W. Maimone, Andrew E. Johnson, Larry M. Scherr

*Jet Propulsion Laboratory, California Institute of Technology,
4800 Oak Grove Drive, Pasadena, CA 91109, USA
Email: {reg.willson, mark.maimone, aej, lmscherr}@jpl.nasa.gov*

ABSTRACT

Machine vision systems often have to work with cameras that become dusty during use. Dust particles produce image artifacts that can affect the performance of a machine vision algorithm. Modeling these artifacts allows us to add them to test images to characterize an algorithm's sensitivity to dust and help develop counter measures. This paper presents an optics-based model that simulates the size and optical density of image artifacts produced by dust particles. For dust particles smaller than the aperture area the image artifact size is determined by the size of the lens aperture and not the size of the particle, while the artifact's optical density is determined by the ratio of the particle and aperture areas. We show how the model has been used to evaluate the effect of dust on two machine vision algorithms used on the 2003 Mars Exploration Rovers.

1. DUST ARTIFACTS IN IMAGES

In the pinhole camera models used in machine vision exactly one ray of light from a point in object space will pass through the camera's pinhole to strike the image plane. With a real lens however, light from a point in object space is collected from a solid angle of rays and projected through the lens onto the image plane, as illustrated in Fig 1. The extent of this solid angle of rays is limited by the lens elements and by the diameter of any diaphragms along the optical path. The limiting diaphragm is called the aperture stop of the lens. The entrance pupil of the lens is the image of the aperture stop as it would be seen if viewed from an axial position in front of the lens.

To model the effect of dust on an image we follow the path of light collected by the lens for a single pixel and consider how dust particles on the lens affect the light reaching the pixel.




Fig. 1. Image formation for a simple lens.

We can call the solid collection angle subtended by a pixel the *collection cone* for the pixel. If a dust particle absorbs or scatters light away from the collection cone (Fig. 2), the light reaching the pixel will be decreased by a factor equal to the fraction of the collection cone blocked by the particle. We call this a *dark dust artifact*.




Fig. 2. Dust blocking a pixel's FOV.

If, on the other hand, a light shining on the lens window is scattered into the collection cone by a dust particle (Fig. 3), then the light reaching the pixel will increase by an additive amount that depends on the intensity of the window illumination. We call this a *bright dust artifact*.




Fig. 3. Dust scattering sunlight into a pixel's FOV.

Bright dust artifacts occur when dust on a lens window is illuminated with an intense light source such as the sun. To help mitigate this lenses are sometimes fitted with sun shades to reduce the range of angles the front lens surface can be illuminated from.

Since absorbing or scattering light away from a pixel's collection cone is much easier to do than scattering light into the collection cone, this paper focuses on dark dust artifacts.

We note that this is a purely geometric model. For visible light and the 10 micron and larger particle sizes we are concerned with scattering follows the laws of geometric optics.

2. DUST MODEL

We represent dark dust artifacts as an array with an attenuation factor for each pixel in the camera. Dark dust artifacts are incorporated into a given test image by simply multiplying the test image by the attenuation image.

To create the attenuation image we find where each pixel's collection cone crosses the dusty lens surface and then compute the fraction of the collection cone blocked by dust particles. Fig. 4 shows the first-order optics model used for these calculations. The optical parameters for the model are

- a diameter of the entrance pupil of the lens
- f focal length of the lens

- s object distance
- s' image distance
- $f/\#$ relative aperture of the lens
- c diameter of the entrance pupil on the window
- w distance between window and the first principal point or front nodal point of the lens
- x_i, y_i position of the pixel on the focal plane array
- x_w, y_w position of entrance pupil image on the window

The relationship between s , s' and f is given by

$$\frac{1}{s} + \frac{1}{s'} = \frac{1}{f} \quad (1)$$

The relative aperture for a lens is given by

$$f/\# = \frac{f}{a} \quad (2)$$

The diameter of the entrance pupil on the window can be approximated by

$$c = a \left(1 - \frac{w}{s} \right) \quad (3)$$

The dust particle parameters are:

- d diameter of the dust particle
- x_d, y_d position of the dust particle on the window




Fig. 4. First-order optical model for dust artifacts.

2.1 Attenuation Image Generation

To generate dust artifacts for an image we need the values of f , $f/\#$, w , and s or s' for the camera system. The first two parameters are typically provided with the camera system, while w can be directly measured or estimated. For the final parameter we can use the image distance s' which is equivalent to the effective focal length of the pin hole camera model produced during the calibration of a machine vision camera model for the camera system.

Given a list of dust particles with diameter d and location (x_d, y_d) on the window, their attenuation image can be generated as follows:

For each pixel position (x_i, y_i) in the image we first determine the location of the intersection of the pixel's line of sight (x_w, y_w) with the window using similar triangles

$$x_w = \frac{w}{s'} x_i \quad y_w = \frac{w}{s'} y_i \quad (4) \quad (5)$$

For each dust particle we then calculate an attenuation factor by determining the area of overlap between the dust particle and the pixel's collection cone. For simplicity we model both the dust particles and the intersection of the pixel's collection cone with the window as circular disks. The attenuation factor is then the area of overlap between two circles of diameter c and d with centers (x_w, y_w) and (x_d, y_d) , divided by the area of the collection cone's intersection with the window, $\pi c^2/4$.

Assuming the dust particles themselves do not overlap, the net attenuation factor for the pixel is one minus the sum of the attenuation factors for the individual dust particles.

Fig. 5 shows how the collection cones for three pixels interact with a dust particle to produce varying attenuation (or transmission) across a detector. The purple cone has zero attenuation because it does not pass through the dust particle. The blue cone has moderate attenuation because it only passes through part of the dust particle. The yellow cone has maximal attenuation because the dust particle is contained entirely within the cone.




Fig. 5. Collection cones for three pixels interacting with a dust particle.

Fig. 6 shows an attenuation image generated for 10 randomly distributed particles from 0.1 to 0.5 mm diameter for the MER field test descent camera. For comparison, Fig. 7 is a field test descent camera image showing artifacts from dust particles deposited during helicopter takeoff. The contrast in this image has been stretched for better viewing so the dust artifacts appear darker than normal.




Fig. 6. Attenuation image generated for the MER field test descent camera.




Fig. 7. Field test descent camera image showing artifacts from dust particles deposited during helicopter takeoff.

2.2 Artifact Size, Shape and Optical Density

The size, shape, optical density of dark dust artifacts are determined by the relative sizes and shapes of the particle compared to the clear aperture c on the lens surface. Note that particles on a particular lens surface all have the same distance w from the pupil. The closer dust is to the pupil, the more pixels are covered by the artifact shadow. At the pupil (i.e. for $w = 0$), the whole focal plane array is uniformly attenuated.

Geometrically the image artifacts are the convolution of the dust particle's silhouette with the collection cone's cross section at the lens surface.

Artifacts can be considered in three size groups:

1. For particle sizes somewhat smaller than the collection cone diameter, the size of the image artifact is essentially constant depending only on the aperture size, the shape of the artifact is the shape of the aperture, and the maximum attenuation within the artifact is the ratio of the particle's area to the cross sectional area of the collection cone (Fig. 8).




Fig. 8. Artifact profile for a small dust particle.

2. For particle sizes approaching the collection cone diameter, the size of the artifact is related to the collection cone diameter plus the particle's diameter, the shape of the artifact is the convolution of the particle's silhouette and the collection cone's cross section, and the attenuation within the artifact approaches 100% at its center (Fig. 9).




Fig. 9. Artifact profile for a medium dust particle.

3. For particles somewhat larger than the collection cone diameter, the size of the artifact is dominated by the size of the particle, the shape of the artifact is the shape of the particle's silhouette, and the attenuation within the artifact is 100% except in the transition around the edges of the artifact (Fig. 10).




Fig. 10. Artifact profile for a large dust particle.

This paper concentrates on group one.

3. DUST MODEL VALIDATION

To validate our dust model we randomly distributed reference particles ~ 0.25 , ~ 0.5 and ~ 1 mm in diameter over a window mounted to the front of camera lens and then took test images with the camera focused on a uniform white target. Test images were taken at three relative aperture settings: $f/22$, $f/16$, and $f/11$ (Figs. 11, 12 and 13) corresponding to entrance pupil diameters of 1.1, 1.5 and 2.2 mm for the 24 mm focal length lens.

We then generated attenuation images using the camera system's parameters ($f = 24$ mm, $w = 22$ mm, $s = 800$ mm) and a randomly distributed field of 0.25, 0.5 and 1.0 mm diameter particles for the same three aperture settings (Figs. 14, 15 and 16).

The diameters of the dust artifacts in the test images were then compared with those in the attenuation images. The results, listed in Table 1, show good agreement. The variability in real reference particle sizes precluded comparing predicted and measured optical densities, but as expected for group one, the test images show artifacts all about the same size, but with three levels of optical density.

Table 1. True versus modeled dark dust artifact diameters.

Relative aperture	True diameter [pixels]	Modeled diameter [pixels]
$f/22$	90-102	98
$f/16$	133-135	134
$f/11$	191-203	195



Fig. 11. Test image taken at $f/22$.




Fig. 14. Attenuation image for $f/22$.



Fig. 12. Test image taken at $f/16$.




Fig. 15. Attenuation image for $f/16$.




Fig. 13. Test image taken at $f/11$.




Fig. 16. Attenuation image for $f/11$.

4. MER DIMES EXAMPLE

The Mars Exploration Rover (MER) Descent Image Motion Estimation System (DIMES) [1] was developed to estimate horizontal wind velocity from images taken during the lander's descent. DIMES takes three descent images as inputs. Using the lander surface-relative attitude and altitude at image capture DIMES warps each image to the ground plane and then computes horizontal displacements between images using image correlation. Image correlation is applied to two locations in the first and second image and two locations in the second and third images. This approach produces four image-based horizontal-velocity estimates that are compared for consistency to each other and to the acceleration computed independently from velocity differences from the lander's inertial measurement unit. Inconsistent results are reported to the lander's control system as a NULL DIMES answer.

One concern during DIMES's development was that dust particles on the descent camera's lens would cause fixed spots in the descent images (Fig. 7) defeating or spoofing the velocity estimation algorithm. Dust could be deposited on the descent camera lens anytime from launch through heat shield separation and first image capture so there was no a priori way of knowing where dust artifacts might appear in the images. Image contrast within the correlation windows was expected to be as low as 0.5% for some of the proposed landing sites so even small dust particles could be a potential problem.

Several different strategies were considered to mitigate the effect of dust artifacts on the DIMES algorithm. To evaluate these strategies a dark dust artifact model was added to the MOC2DIMES camera simulator used to develop DIMES [2]. The simulator generated attenuation images for 1 to 20 dust particles from 0.1 to 0.25 mm in diameter. The resulting dark dust artifacts were 18 pixels in diameter and from 2.6 to 16% attenuation. The attenuation images were then applied to the Mars Orbital Camera (MOC) images of candidate landing sites used as test images for DIMES. Fig. 17 shows a MOC test image with added dark dust artifacts. The intensity gradient in the image is the result of the camera's optics and the CCD's short exposure time and does not affect the DIMES algorithm. Fig. 18 shows the results for the baseline DIMES algorithm for a total of 100 DIMES simulation runs.

Based on the simulation results it was decided the baseline DIMES algorithm was sufficiently robust to one or two dust particles that additional algorithmic complexity was unwarranted.




Fig. 17. Sample MOC test image with dark dust artifacts (shown before the final 4:1 row binning step).




Fig. 18. Percent NULL DIMES answers for different numbers of dust particles.

5. MER STEREO EXAMPLE

The MER rover's stereo cameras are used to generate range data for hazard detection during driving and provide maps of surface position and orientation for planning tool placement and traverses.

MER's stereo vision software [3] can generate a range estimate for every pixel in an image. Given pixel coordinates in one image, it correlates a small window around that pixel with potential matching windows in the other image. In simple terms, this produces a standard correlation curve; the peak of the curve indicates the "best" match.

However, not every estimate is considered reliable. For example, if no obvious peak is found in the correlation curve, that range value is thrown out. A range estimate is also thrown out if the maximum correlation value lies at the tail end of the correlation curve (implying that the actual peak lies in an un-tested area). Finally, after all other consistency checks have run, a blob filter is applied to prune out even more data, by eliminating unconnected regions smaller than some fixed number of pixels. These consistency checks prune out unreliable range values, which allows the rover to maintain high confidence in those that remain.

During the rover's development there was a concern that dust produced by the Rock Abrasion Tool (RAT) would contaminate the rover's front Hazcams and produce image artifacts that could disrupt the left-right image correlation and cause a loss of range data. Studies suggested that the bulk of the particles generated by the RAT would be 45 microns or smaller.

Since MER stereo Hazcams were unavailable, to evaluate the potential impact of dust on ranging performance dark dust attenuation images were generated and applied to 10 pairs of stereo images of Martian terrain taken with the Mars Pathfinder IMP camera [4]. Fig. 19 shows one Pathfinder image with and without dust artifacts. The stereo range maps produced from the images with the dust artifacts were then compared to the range maps produced from the original dust free images. The metric used was the ratio of range image density from images with dust artifacts to the range image density from the original dust free images. As mentioned above, the stereo vision processing has some built-in consistency checks that cause range data to be thrown out. The noisier the images are, the more range data will be thrown out.




Fig. 19. Pathfinder image with and without simulated dark dust artifacts.

5.1 Test Results

A total of 3780 separate test cases were run on the stereo vision software with dust particle counts from 20-2000, particle diameters from 5-45 microns, stereo pyramid levels from 1-3, and stereo blob filtering from 100-2590 pixels. Each test case produced a range density which was divided by the range density achievable using the original, unmodified images. Even in the worst case dust artifacts decreased the range density by only 3%. Thus the Pathfinder IMP image tests suggested the presence of 2000 or fewer particles of size 45 microns or smaller would have a negligible effect on MER's Hazcam stereo ranging performance.

5.2 MER Flight Experience

Over the course of operations both Spirit and Opportunity's Hazcams have experienced dust contamination without noticeable degradation in their stereo ranging performance. Figs. 20 and 21 show images from Spirit's right front Hazcam before and after a dust contamination event around Sol 417 as it was using the RAT on a rock in the camera foreground. The image artifacts are consistent with numerous small particles deposited across the lens. Opportunity's rear Hazcams experienced similar dust contamination around Sol 331 as it was driving around the heatshield site. Figs. 22 and 23 show images from Opportunity's right front Hazcam before and after a dust contamination event around Sol 471 as it was driving out of a sand dune. The artifact visible against the sky is ~17 pixels across and has an optical density of ~77% which would correspond to a circular dust particle at the front window (25 mm from the pupil), ~90 microns in diameter.

6. CONCLUSIONS

Models of the image artifacts produced by dust on camera optics can be invaluable for evaluating the impact of dust contamination on machine vision systems. Using only four parameters the proposed optics-based dust artifact model produces realistic image artifacts for arbitrary sizes and distributions of dust particles. While the results presented focused on dark artifacts produced by particles smaller than the entrance pupil diameter, the principles can be extended to bright artifacts caused by illuminated dust and to arbitrarily large particles.

7. ACKNOWLEDGMENTS

The research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

8. REFERENCES

1. Cheng Y., Johnson A.E., and Matthies L.H., *MER-DIMES: A Planetary Landing Application of Computer Vision*, IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2005), June 2005.
2. Willson R.G., Johnson A.E., and Goguen J.D., *MOC2DIMES: A Camera Simulator for the Mars Exploration Rover Descent Image Motion Estimation System*, 8th International Symposium on Artificial Intelligence, Robotics and Automation in Space, September 2005.
3. Goldberg S.B, Maimone M.W., and Matthies L.H., *Stereo Vision and Rover Navigation Software for Planetary Exploration*, IEEE Aerospace conference proceedings, volume 5, Big Sky, Montana, USA, pp 2025-2036, March 2002.
4. Smith P.H., et al., *Results from the Mars Pathfinder camera*, Science, Vol 278, Issue 5344, pp 1758-1765, December 1997.




Fig. 20. Spirit right front Hazcam on Sol 416 before dust contamination.




Fig. 22. Opportunity right front Hazcam on Sol 470 before dust contamination.




Fig. 21. Spirit right front Hazcam on Sol 417 after dust contamination.



Fig. 23. Opportunity right front Hazcam on Sol 471 after dust contamination.