AIAA JOURNAL
Vol. 50, No. 5, May 2012

# Texture-Based Photogrammetry Accuracy on Curved Surfaces

Alan Jennings* and Jonathan Black†
*U.S. Air Force Institute of Technology, Wright-Patterson Air Force Base, Ohio 45433*

DOI: 10.2514/1.J050956

**Photogrammetry is a preferred technique for noncontact measurement of flexible structures, such as large membranes and flapping-wing vehicles. Traditionally, discrete features are triangulated between photos. New techniques use the local texture of a surface to match random speckle among images. Texture-based photogrammetry allows for higher resolution surfaces, and this paper tests if the precision is similar to traditional photogrammetry. Test surfaces provide known, rounded profiles used to compare surface reconstructions. Many images are taken of an object with a projected texture, and the depth error is used to quantify the accuracy of the results. Results on surface meshes from 800 to 2500 points show accuracy on the order of 1:2000, or nearly that of one pixel. Camera locations had a surprisingly minor effect on surface quality. There is a slight correlation of more parallel views to more points on the surface. To show that the technique extends to practical use, results are shown for a flapping cycle of a membrane wing. Dense surface meshes are important for higher fidelity models in computational fluid dynamics and finite element analysis.**



## I. Introduction

FOR a measurement system to be effective, it must be able to observe the normal behavior of an object without altering the behavior significantly. Large flexible space structures and micro air vehicles are challenging to measure because of their large range of motion during natural operation. In addition, mounting sensors on them would add significant inertia. Traditional assumptions lose accuracy with novel materials and very thin structures, requiring the validation of new models [1]. Traditionally, wings were modeled as rigid, but wing flex is sometimes designed into the system [2]. Many ornithopters use flexible membranes, despite the lack of established theory, promoting a more empirical design strategy [3–6]. Membrane dynamics have also been of interest to the gossamer spacecraft and weather balloon communities to understand large membrane deployment. Recent interest in membrane dynamics has grown based on the sun shields for the James Webb Space Telescope (JWST) [7]. The JWST requires well-modeled membrane dynamics to ensure accurate optics and proper heat rejection.

The two most popular noncontact measurement methods use lasers or images. Lasers measure a single point at a time, so many trials are required to capture surface shapes, whereas photogrammetry has a full-field view and can capture the surface's shape and motion from a single test [8,9]. Time-of-flight laser scanning is commonly used on large static scenes, such as buildings. Laser triangulation is commonly used for small static scenes, such as inspecting a manufactured part. Geometric validation of the JWST was conducted using novel interferometry techniques [10]. Capacitance measurements and full-field interferometry using moiré patterns give good results but require precise alignment, a mostly flat surface and a conducting surface in the case of capacitance [8,11]. Recently, a hawk moth wing was captured in-flight via a moiré pattern [12]. Laser scanning has been fused with feature-based photogrammetry to capture large and small features [13]. Other recent work incorporates tracking an object in real-time with a vibrometer for vibration results at one location or multiple cameras for surface results [14,15].

A good introduction to photogrammetry is given by Gruen [16]. With digital imaging and sophisticated software techniques, photogrammetry has developed subpixel resolution [17,18], practical accuracy [19,20], and processing speed. Example uses include as-built computer aided drafting [13,20], airborne surveying [21], and forensic reconstruction. A limitation is the often labor-intensive process of feature matching despite advances in automatic feature extraction [13,22]. Methods have been developed to accentuate features, many by adding or projecting high contrast targets onto the three-dimensional (3-D) surface [9,22]. Rather than adding targets, texture-based photogrammetry uses surface texture to create features that are matched by texture comparison. In this manner, the number of targets does not limit the surface resolution and allows for the creation of very high density point profiles.

Videogrammetry offers the ability to use a sequence of synchronized images from multiple viewpoints to capture surface shape as it changes. Wind-tunnel videogrammetry results for 3-D motion of a carbon fiber composite miniature unmanned aerial vehicle concept are presented in Black et al. [23]. These results identify resonant frequencies and mode shapes as a function of wind speed and angle of attack. Dynamics of a membrane have also been measured via videogrammetry [24].

New texture-based photogrammetry techniques have the ability to produce dense surfaces consisting of thousands of points. These surfaces can be used for finite element analysis and computational fluid dynamics (CFD) analysis [25,26]. Such texture-surface reconstruction techniques using synchronized videos are the most recent adaptation of this approach for dynamic characterization of various lightweight and inflatable space structure test articles [27–35]. This paper studies how the expected accuracy of texture photogrammetry may be influenced by camera locations and shows texture photogrammetry feasibility on dynamic surfaces.

To have confidence in a measurement system's results, its accuracy must be quantified. Unfortunately, the accuracy of photogrammetry is very hard to quantify. Accuracy estimates for photogrammetry are developed by software vendors but do not take factors, such as feature definition, into consideration. Heuristics, such as a set of best practices, are often used to design the layout of cameras with good results [20], but so-called poor layouts have also given good results and vice versa. Previously, target photogrammetry has been validated by laser scans [19,20,36] and capacitance [8]. Work has also been done to compare different photogrammetry techniques [36].

By testing a variety of known surfaces from different views, the dependence of measurement error on physical setup can be investigated to empirically develop a set of best practices for texture-based

Presented as Paper 2010-2749 at the 11th AIAA Gossamer Systems Forum, Orlando, FL, 12–15 April 2010; received 24 September 2010; revision received 27 October 2011; accepted for publication 24 November 2011. 

*Research Assistant, Department of Aeronautics and Astronautics, 2950 Hobson Way. Student Member AIAA.
†Associate Professor; jonathan.black@afit.edu. Associate Fellow AIAA (Corresponding Author).

photogrammetry. The results presented also validate accuracy on known, curved surfaces and show representative results for large membrane deflection. In the next section, the quantitative measures will be described. The equipment used, along with the test objects, will be described in Section III, and the theory specific to texture-based photogrammetry is given in Section IV. Section V presents the results for different surface types showing accuracy, point distributions, and surface effects. Results are also qualitatively described identifying types and causes of poor quality. To show that texture-based photogrammetry can be applied for practical purposes, Section VI demonstrates results from a flapping membrane wing that could later be used in a CFD analysis. The paper concludes with a summary of the major findings in Section VII.

## II. Assessment

The goal of this paper is to systematically investigate the relationship between camera setup and reconstructed surface quality. The setup is characterized by the location and orientation of the camera positions relative to the object. Dependence of accuracy on texture-based photogrammetry algorithm parameters, such as matching region size and measures of texture regularity, are not investigated in this study. Images sets are synchronized, so the motion of surface points is not considered in this investigation. The surface fit quality is quantified by accuracy as measured in the surface depth. Qualitative analysis considers the distribution of points and error in the point cloud.

Error is defined as the height of the point given by photogrammetry minus the surface's known height, estimated using the geometric equation that models the actual surface contour. Traditionally, root-mean-squared (RMS) error has been used to average error, but this is a poor choice due to the high likelihood of outlier points. Outlier points are typical for point cloud methods where thousands of points are generated. Methods for detecting outliers are well established, but the methods are normally application-specific or require tuning. To prevent bias based on outlier detection method, outliers were not removed. The RMS error is a special case with $n = 2$ of the $n$-mean-norm error given by Eq (1):

$$E_n = \frac{1}{N^{1/n}} \left( \sum_{i=1}^{N} |e_i|^n \right)^{1/n} \tag{1}$$

where $e_i$ is the error of point $i$, and $N$ is the number of points. Because the mean in the RMS is taken of the square, the effect of individual points on the mean grows as they move further away from the consensus. When composing a surface representation, overall consistency with the majority of points is more important than outlier points, which are likely wrong. The effect of different values of $n$ in Eq. (1) can be seen in Fig. 1, where the median is also included for comparison. The error from a surface reconstruction is used in Fig. 1 and is a typical distribution. The weighting factor for the RMS grows rapidly, resulting in a shift of the norm towards outlying points. Larger values also affect the 1-mean-norm (the mean) but not at a greater rate than small values. The median is determined only by the ordering of the points, so outliers have no effect based on how far they are from the consensus. Unfortunately, this lack of sensitivity to slight variations fails to provide a search direction for the optimizations used to register surfaces to the truth surface. The $1/2$ and $1/5$ mean-norm are both near the median, but the $1/5$ norm is so evenly weighted, it does not have sufficient sensitivity to effectively distinguish goodness. The $1/2$ norm, shown in Eq. (2), is used to define a good registration and an average error of the surface:

$$E = \frac{1}{N^2} \left( \sum_{i=1}^{N} \sqrt{|e_i|} \right)^2 \tag{2}$$

## III. Systems

### A. Camera

Pictures were taken on a single lens reflex camera in aperture-controlled mode with zoom and focus held constant. Both the aperture and the lens geometry affect camera distortion. To make precise measurements, these distortion parameters must be solved independently (based on calibration images, not test images) and then held constant through all images [37]. A tripod and remote control were used to stabilize the camera during capture so motion blur was not a factor. Files were saved as lossless, color, JPEG images at 2560 by 1920 pixels (5 megapixel) with a nominal file size of 3.5 MB. Shutter time, which does not affect image distortion, was automatically extended as needed for adequate contrast.

### B. Photogrammetry Software

A commercial, general-purpose photogrammetry software was used for performing the surface reconstruction. The process used is presented in Table 1. The process begins with calibrating the camera by determining the parameters of lens distortion. To ensure calibration parameters are not spurious results of optimization, they are checked for correlation to ensure statistical significance. Idealization removes lens distortion (radial distortion with second- and fourth-order terms and lens decentering), centers the photo about the principle point, and makes each pixel square. The photo is padded with zeros to accommodate the entire original image, resulting in a black border. An example idealized photo is shown in Fig. 2.

The photos are automatically prepared and oriented based on coded targets. Coded targets have an outer code ring that is identified by the software. When the same identifying number is seen on multiple photos, a consistency constraint is created to allow for solving for the camera stations. The inner circle is used as a target for high accuracy marking using a least-square matching (LSM) algorithm. LSM uses a template to fit an oval to the target mark to subpixel resolution. The extents of the test surface are marked by

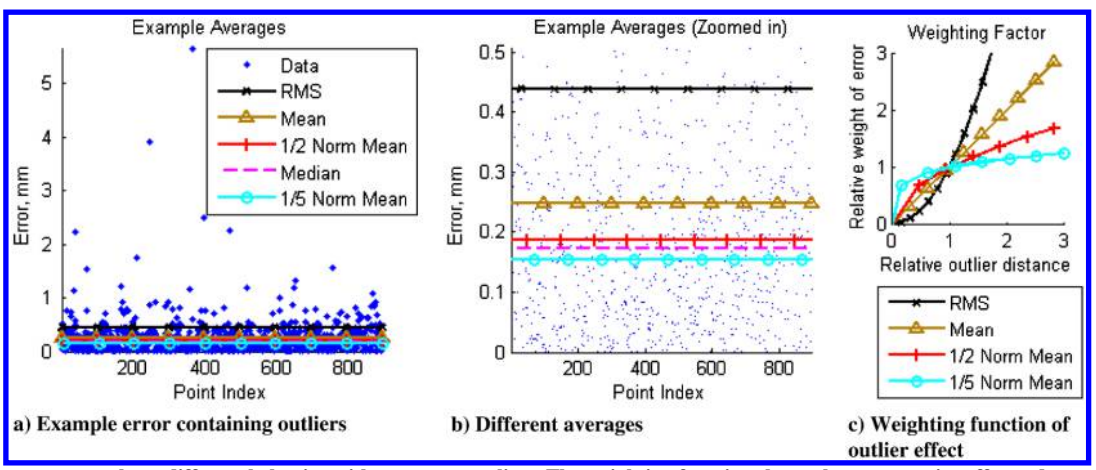


**Fig. 1 Different averages have different behavior with respect to outliers. The weighting function shows the comparative effect a data point has on the result based on how far the point is from the average. The weight increases based on the distance from zero by the power of the norm used.**

**Table 1 Surface mapping analysis process**

1. Take camera calibration pictures at the focus, zoom, and aperture to be used.
2. Place a spread of coded targets around the surface and project texture onto the surface.
3. Take pictures of the surface from various angles.
4. Process camera calibration to determine focal length and lens distortion parameters.
5. Idealize photos to create equivalent pinhole camera images.
6. Using automatic matching of coded targets, determine the camera stations of each image.
7. Manually select corners defining the nominal object boundaries.
8. Define scale and coordinate system from nominal object boundaries.
9. Create a dense surface point cloud from each pair of photos using texture-based photogrammetry.
10. Export point clouds as 3-D points for analysis.
11. Register each mesh to truth surface, calculate average error, and plot relationships of setup to error.

hand, but the consistency constraint for error minimization is not applied, so that the low-precision marks will not affect the high accuracy project. Using the surface formed by the extents as a reference plane, the dense surface mesh tool applies texture-based photogrammetry to create a high-resolution point cloud for each photo pair (see Section IV). These point clouds are then exported and analyzed.

### C. Test Articles

Test articles are made by a 3-D printer to 600 dots per inch (dpi) (18 $\mu$m, 0.0017 in) accuracy horizontally and 1600 dpi (15 $\mu$m, 0.0006 in) in depth to a nominal size of 25–50 mm × 114 mm × 38 mm (1 to 2 × 4.5 × 1.5 in). Shapes were drafted as a solid model using spline interpolation and lofted across the width for 3-D variation. The known surface was verified by a laser scan to ensure interpolation consistency. The collection of objects is shown in Fig. 3. A gentle sine extending the length of the block, Fig. 3a, represents typical curvature. A set of planes, Fig. 3b, were used to find if accuracy increased for planar features. A protruding sine, Fig. 3c, has a peak-to-peak amplitude of half the period. This results in a challenging surface due to high curvature and a wide variety of surface directions with respect to the view angle. A coat of white paint provided a background for projection of a speckle pattern. An LCD projector was used to provide an identical speckle for all test articles. To get a high contrast, the projector was the primary source of light. Figure 2 shows a close up of the protruding sine object with the texture and coded targets.

### D. Speckle Pattern

The proposed texture was designed to replicate the traditional sprayed speckle. It consisted of high-contrast elliptical dots, randomly distributed and projected onto the surface with a digital LCD projector. Each dot had six parameters: $x$ and $y$ locations, local $x$ and $y$ radii, local rotation, and intensity values. Note that images have an inherent length unit, so a baseline length is used to scale other length parameters. The square root of the average background area per speckle ($\sqrt{\text{total area}/\text{number of speckles}}$) was used as the baseline length. Each speckle is made elliptical by distributing the local $x$ radius across [0.25,080], whereas the local $y$ radius was distributed across [0.51,1.13]. The speckle was then rotated by an angle of range $[0, \pi]$ and given an intensity of [0.3,0.5]. All random variables were uniformly distributed across their range, except for radii, which were distributed uniformly in the square, giving a slight preference to larger speckles. With these parameters, there was significant overlap, resulting in a greater diversity of shapes than eclipses, as shown by an example in Fig. 4. To provide a brighter image, the image was inverted (negated) to give a white background and raised two units. Values of 1 or greater were displayed as white, values of 0 or less were black, and intermediate values were shades of gray.

## IV. Texture Photogrammetry Theory

Texture-based photogrammetry relies on the same principles as target or feature-based photogrammetry but uses the randomly varying intensity of a small set of pixels as a target, which is matched to a similar set on another photo. The size of the pixel windows should be large enough to capture descriptive information about the set but small enough to provide distinct statistics and correlations. Texture is not a precisely defined term, yet it is easily qualitatively understood. An excellent source that reviews the concept of texture is

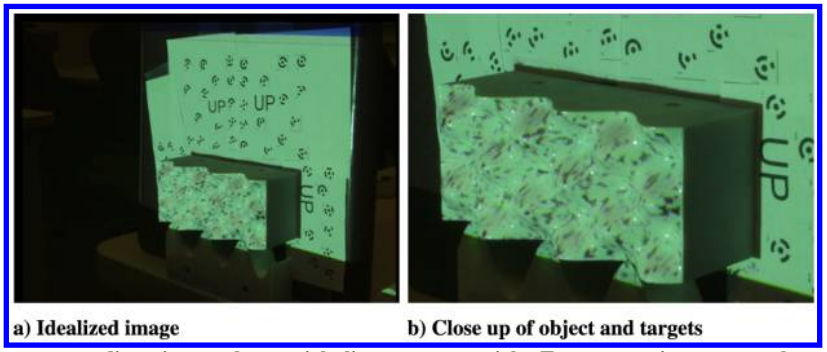


**Fig. 2 The photo idealization removes distortions so that straight lines appear straight. Targets are given a rotary bar code used to uniquely identify them for automatic image orientation. The test object has a texture applied as shown.**

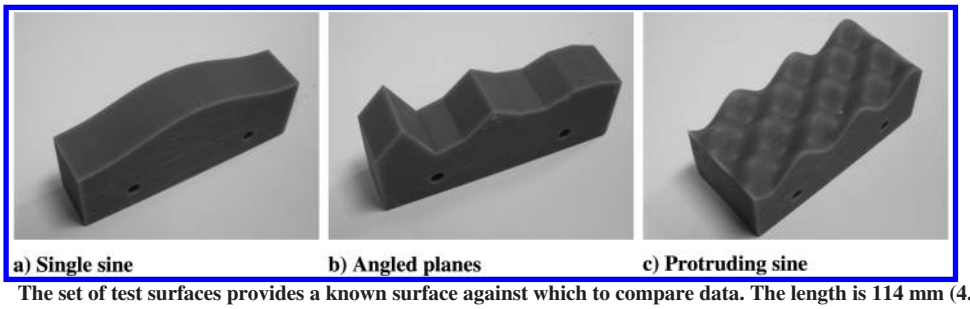


**Fig. 3 The set of test surfaces provides a known surface against which to compare data. The length is 114 mm (4.5 in).**

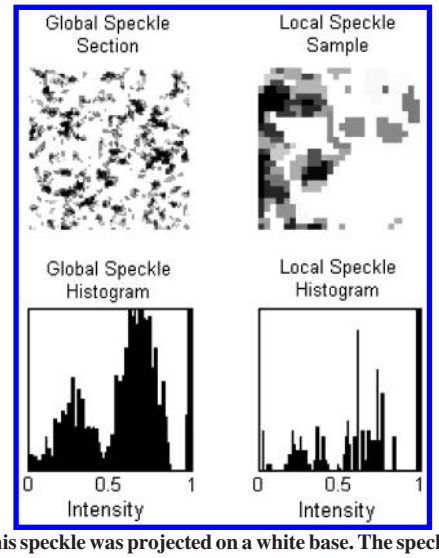


**Fig. 4 This speckle was projected on a white base. The speckles provide a reasonably even distribution of intensity but, locally, have a distinct histogram profile and speckle profile.**

Gonzalez and Woods [38]. Most texture descriptions are based on the distribution of intensity values, as shown in Fig. 5. Many statistical measures exist for characterizing the distribution, with some providing greater sensitivity to certain effects than others. For example, the mean value is very sensitive to lighting and is not a good choice; but second, third and fourth moments of the mean provide information of the spread, skewness, and uniformity of the histogram and are useful for automatic matching. High-order moments of the intensity can be scaled to provide measures invariant to translation, rotation, and pixel scaling. Other techniques, such as correlation, involve shifting the set and observing the similarity but are sensitive to perspectives or other rotations. Shifting techniques use the natural texture as a template and observe how similar the set is to its neighbors. One or more of these comparison techniques can be used to match sets of pixels (targets) in multiple images.

To show how textures measures could be used for matching, measures are shown for the same four points on two images in Fig. 5. The first three rows show measures from the first image, with the corresponding measures from the second image are given below in the last three rows. The first row shows the original images with the local intensity scaled to a 0 to 1 range on the right to enhance contrast. This scaling can remove effects, such as a brighter photo, but can sometimes accentuate differences, as is the case with point 3. The histogram shows how often an intensity is found in the images and is used to identify multiple peaks or other distribution shapes. The final row shows a set of statical measures that compare the moment of the intensity normalized to be constant despite translation, rotation or scaling. The first two points have good comparison, whereas the fourth only matches well in the histogram. The identical points were selected by hand and only scaled. The roughly 90 deg rotation between photo 1 and 2 is still seen, showing that perspective has not been addressed. For a high-fidelity match, perspective and cross-correlation should also be used.

A pixel from one image matched to a pixel of another image correspond to a 3-D location based on the orientations and positions of the cameras. A set of pixels and their matches correspond to a 3-D surface. Because camera positions and orientations are known, a location in the image plane (determined by pixel location) has a distinct ray through space extending from the focal point through that location in the image plane. Rays for the same surface point in two photos will intersect at the 3-D location of the point.

If the angle between the two rays is small, then slight changes in either ray would produce large differences in point depth [16]. Because this sensitivity to ray uncertainty is undesired, larger ray angles are desired. One solution is to have the camera viewing axes at a broad angle, the maximum being 90 deg, so deviations corresponding to image pixels would still have broad angles. This approach is recommended for target-based photogrammetry [20].

However, textures can be more distorted at wide viewing angles due to reflections and perspective changes. For better correlation, standard perspective transforms can be used to reshape textures from one photo into the other photo's view or to a compromise view between both cameras. Compensating for changes in appearance due to reflection or other incident effects requires specific knowledge

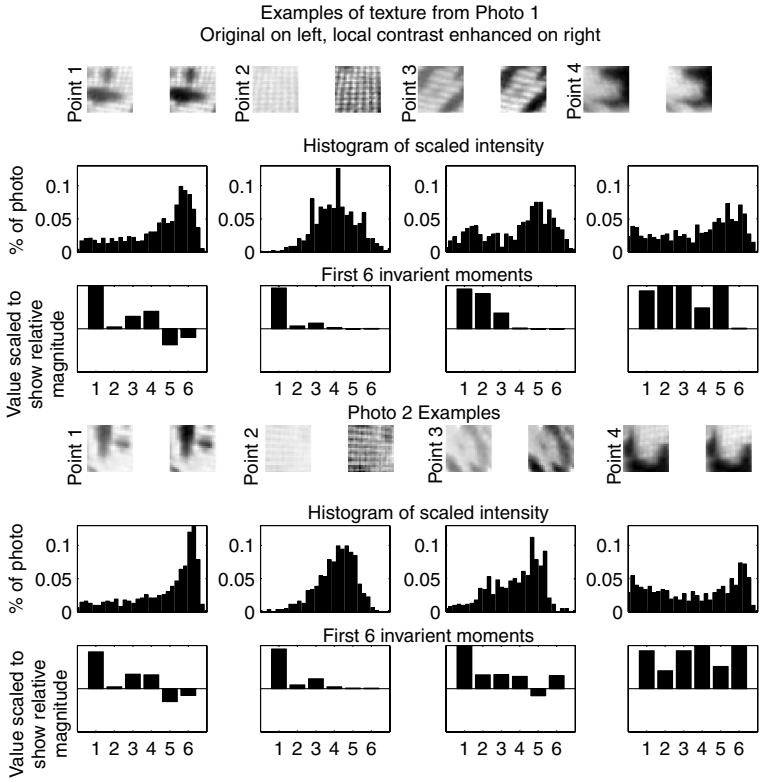


**Fig. 5 Intensity distributions for four points for texture-based matching.**

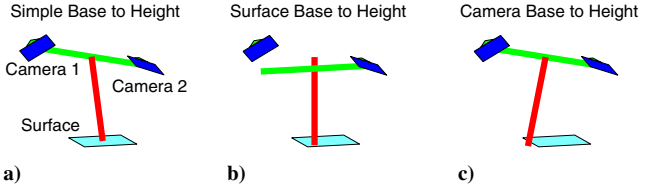


**Fig. 6 The base-to-height ratio compares the distance between cameras to their distance to the surface. A) A simple ratio does not consider the orientation of the vectors. B) The surface and c) camera ratios put the distances in terms of their perpendicular directions.**

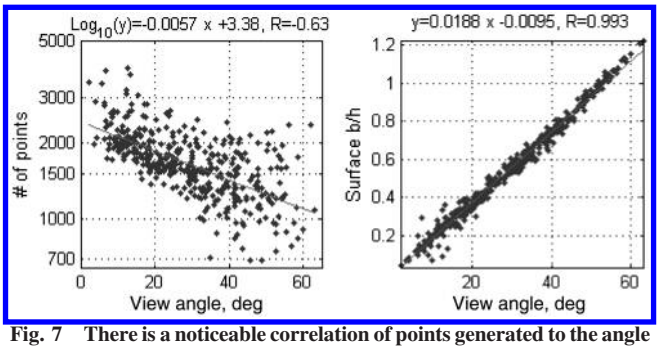


**Fig. 7 There is a noticeable correlation of points generated to the angle of view separation but not a conclusive relation. This means that where there is a small angle of separation, there tends to be more points, although the correlation coefficient only has a magnitude of 0.63. This correlation also extends to the base-to-height ratio because the two values are closely correlated.**

about the type of surface and is difficult, if not impossible, to accomplish reliably. For these reasons, texture-based photogrammetry recommends a set of near-parallel viewing axes to reduce error introduced by view direction [37].

The two rays through space result from the separation of the cameras, referred to as the baseline. A small baseline results in triangulation sensitivity, whereas a large baseline results in two very different perspectives due to viewing angle and reflection. A baseline of $1/10$ to $1/2$ the distance to the surface is recommended [37]. The disadvantage of moving the cameras further apart is that the object will be covered by fewer pixels, increasing pixelation. Nominally this would produce a point cloud with fewer points; but because sets of pixels are used, superresolution techniques [17,18] can be used to obtain a commensurate number of points as less pixelated images but with greater uncertainty.

Two terms that have arisen to characterize a setup are the camera angle and the base-to- height ratio. Unfortunately, there are a number of valid definitions of these terms. The software used here measures camera angle by the angle between the two principle axes of the images used. Measuring the included angle between the two rays from the surface center to the image seems more appropriate. These rays reflect the actual viewing angle of the surface on the image, rather then the image in general. For example, panning the camera will change the camera angle; but as the surface moves across the image, the actual rays through space stay the same. For the base-to-height ratio, the simplest approach uses the distance between camera stations as the base and the average distance from the surface center to the camera stations as the height. The result of this measure is that staggered pairs (both images on the same principle axis with one spaced well in front of the other) have a good base-to-height ratio but offer poor geometry. Two alternative methods find the perpendicular component of the other vector, as contrasted to the simple base-to-height ratio in Fig. 6. A surface base-to-height ratio considers only the height component perpendicular to the base and only the base component parallel to the surface. A camera base-to-height ratio considers only the height component perpendicular to the line between cameras. These produce slight increases in predictive

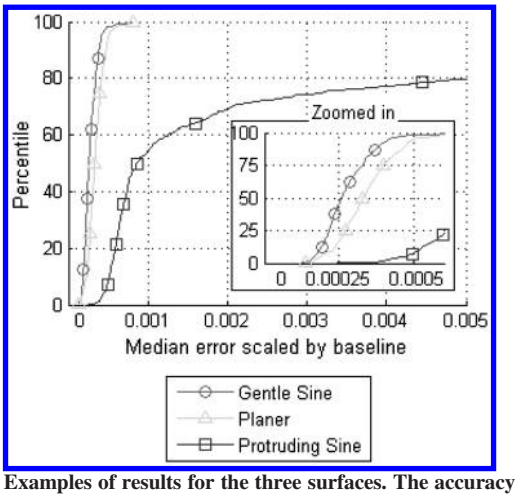


**Fig. 8 Examples of results for the three surfaces. The accuracy is the inverse of the normalized error.**

capability, but one is not significantly better than the other. If the surface fills most of the image, as is the case when the surface has the maximum resolution, then the surface referenced base-to-height ratio would be related to the viewing angle.

In summary, there are a set of competing factors that can heuristically be applied but are difficult to solve for optimality. The guideline is that the object should be two to 10 times as far away from the cameras as the distance between cameras with the object filling the largest possible area while still getting sharp photos [37]. Some methods have been used to find the best setup, such as genetic

**Table 2 Accuracy distributions for gentle sine and angular planes**

| | | | Base-to-height ratio | | | | |
|---|---|---|---|---|---|---|---|
| | | | 0.0 → 0.1 | 0.1 → 0.3 | 0.3 → 0.5 | 0.5 → 1.0 | >1.0 |
| Camera angle | 0 deg → 15 deg | # of meshes | 14 | 120 | 38 | 4 | |
| | | 25th percentile | 1:2, 000 | 1:2, 200 | 1:2, 400 | | |
| | | 50th percentile | 1:2, 300 | 1:2, 700 | 1:3, 000 | 1:3, 800 | |
| | | 75th percentile | 1:2, 400 | 1:3, 100 | 1:3, 700 | | |
| | 15 deg → 30 deg | # of meshes | 2 | 42 | 154 | 160 | |
| | | 25th percentile | | 1:2, 500 | 1:2, 500 | 1:2, 900 | |
| | | 50th percentile | 1:3, 000 | 1:2, 900 | 1:3, 000 | 1:3, 400 | |
| | | 75th percentile | | 1:3, 600 | 1:3, 800 | 1:4, 100 | |
| | 30 deg → 45 deg | # of meshes | | | 2 | 194 | 2 |
| | | 25th percentile | | | | 1:3, 400 | |
| | | 50th percentile | | | 1:4, 200 | 1:3, 800 | 1:5, 300 |
| | | 75th percentile | | | | 1:4, 500 | |
| | >45 deg | # of meshes | | | | 32 | 40 |
| | | 25th percentile | | | | 1:4, 000 | 1:4, 300 |
| | | 50th percentile | | | | 1:4, 600 | 1:5, 100 |
| | | 75th percentile | | | | 1:5, 500 | 1:5, 900 |

**Table 3 Accuracy distributions for protruding sine**

| | | | Base-to-height ratio | | | | |
|---|---|---|---|---|---|---|---|
| | | | 0.0 → 0.1 | 0.1 → 0.3 | 0.3 → 0.5 | 0.5 → 1.0 | >1.0 |
| Camera angle | 0 deg → 15 deg | # of meshes | 10 | 98 | 16 | | |
| | | 25th percentile | 1:500 | 1:100 | 1:600 | | |
| | | 50th percentile | 1:1, 100 | 1:600 | 1:1, 100 | | |
| | | 75th percentile | 1:1, 300 | 1:1, 100 | 1:1, 600 | | |
| | 15 deg → 30 deg | # of meshes | 3 | 36 | 110 | 54 | |
| | | 25th percentile | 1:900 | 1:800 | 1:200 | 1:300 | |
| | | 50th percentile | 1:1, 200 | 1:1, 400 | 1:1, 000 | 1:1, 100 | |
| | | 75th percentile | 1:1, 900 | 1:1, 600 | 1:1, 500 | 1:1, 500 | |
| | 30 deg → 45 deg | # of meshes | | | 9 | 155 | 1 |
| | | 25th percentile | | | 1:800 | 1:300 | |
| | | 50th percentile | | | 1:1, 700 | 1:1, 300 | 1:1, 300 |
| | | 75th percentile | | | 1:2, 100 | 1:1, 700 | |
| | >45 deg | # of meshes | | | | 49 | 160 |
| | | 25th percentile | | | | 1:800 | 1:300 |
| | | 50th percentile | | | | 1:1, 400 | 1:1, 300 |
| | | 75th percentile | | | | 1:1, 700 | 1:1, 700 |



algorithms, and validate photogrammetry best practices, whereas others suggest nontraditional camera layouts [39]. These works are often focused on a few types of uncertainty, such as ray and pixelation sensitivity, but often use idealized targets unlike texture techniques. This work seeks to measure practical accuracy on nonuniform surfaces whether or not conventions are followed.

## V. Results

The procedure in Table 1 was performed on the objects in Fig. 3 with the texture in Fig. 4 projected on their faces as shown in Fig. 2. Sets of images used stereo views up to separation angles of 60 deg and base-to-height ratios of 0 to 1.2. Figure 7 shows that the number of points generated for a mesh is related to the camera angle and the base-to-height ratio. Using the maximum angle of the camera principle axis to the surface normal did not produce a correlation. The different base-to-height measures and camera angles have strong correlations, as shown by the linear trend in the right hand plot of Fig. 7, making it impossible to separate effects of one from the other. The baseline to scale the error is the distance across the image diagonal in the plane of the surface (450 mm compared with a length of 120 mm across the test object). Because each pair of 20–40 photos was used, hundreds of meshes were created. The camera locations were determined using target-based photogrammetry with all of the images. This gives the strongest estimate of camera locations and separates target-based photogrammetry error from texture-based photogrammetry error. Results show that texture-based photogrammetry is very robust, assuming cameras positions are accurately determined, because it yields good results even if best camera spacing practices are not followed.

Surface generation techniques have high relative accuracy, meaning that the relative position between points is very accurate, but there could be systemic error in the form of translation or rotation of the points together. The accepted practice is to use a registration algorithm to align multiple sets of data. A simple registration was used by minimizing the $1/2$ norm error to the truth surface by nonlinear programming over parameters for translation, rotation, and a scale factor. Though the optimization uses the truth surface, which will not be available in a test scenario, it will not provide spurious enhancements due to the large ratio of points in the data sets to optimization parameters (thousands of data points to seven optimization parameters). The adjustments were minor, typically less

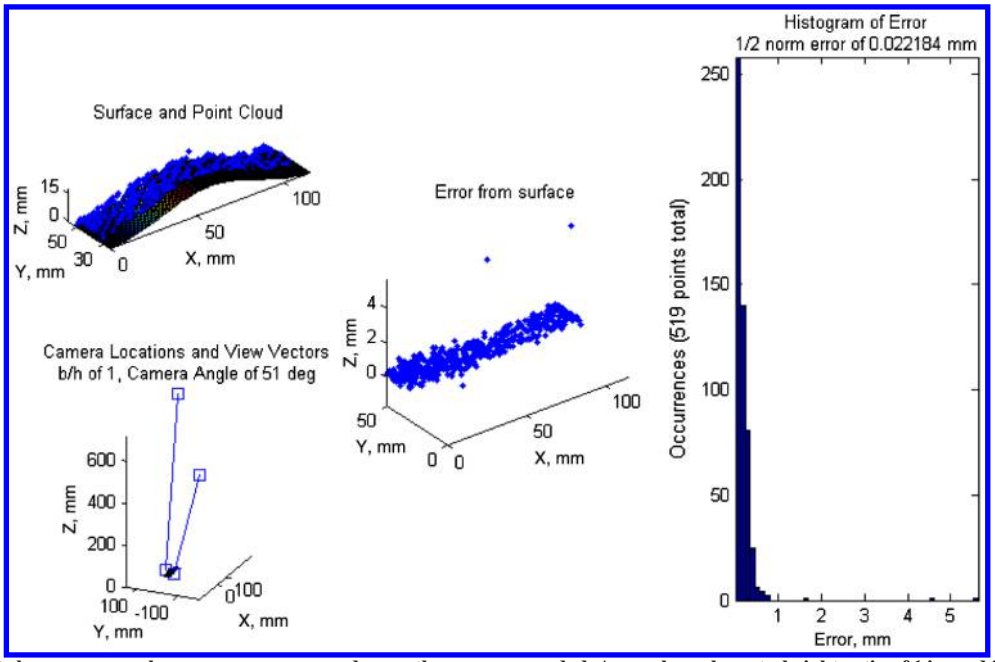


**Fig. 9 Good results happen even when cameras are spread more than recommended. A very large base-to-height ratio of 1 is used for these results. On the left, results are shown over the truth surface with the relative camera locations shown in the lower plot. The center plot shows the error of the result by location, whereas the right hand plot is a histogram of error occurrence.**

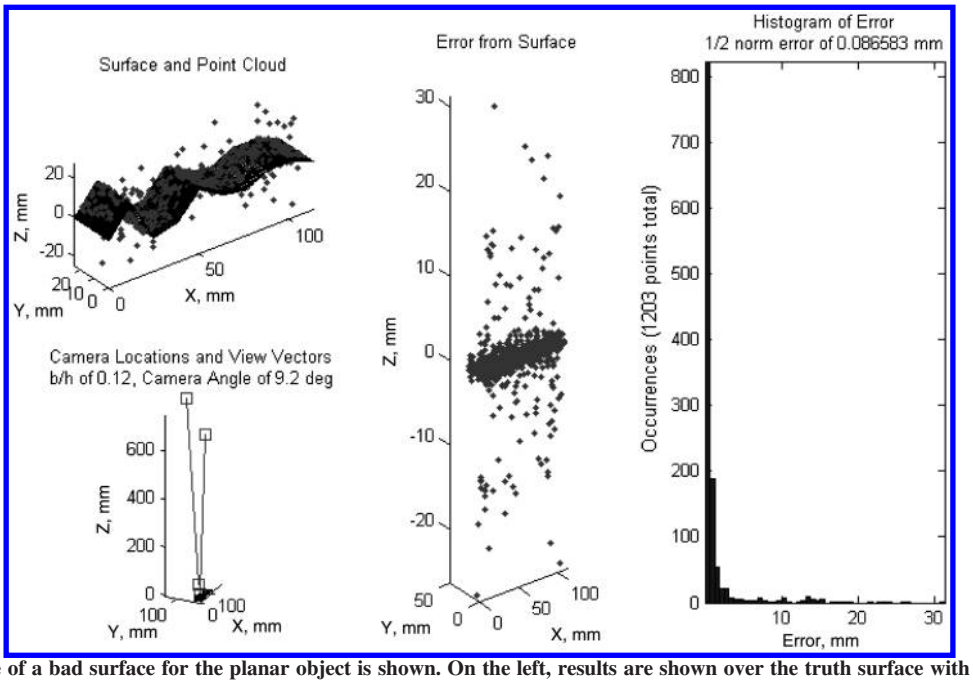


**Fig. 10 An example of a bad surface for the planar object is shown. On the left, results are shown over the truth surface with the relative camera locations shown in the lower plot. The center plot shows the error of the result by location, whereas the right-hand plot is a histogram of error occurrence. Note that though there are many errant points, there are many more accurate points resulting in a low overall error. The base-to-height ratio is near the lower limit of photogrammetry guidelines.**

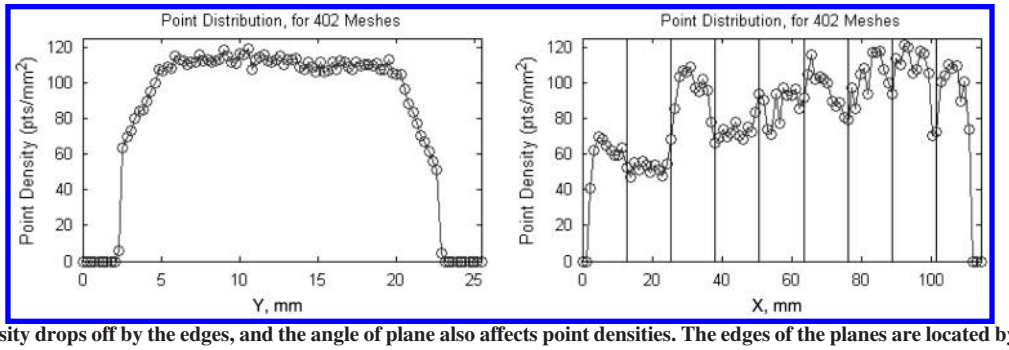


**Fig. 11 Point density drops off by the edges, and the angle of plane also affects point densities. The edges of the planes are located by the vertical lines. The two steepest angles are on the left and have densities slightly greater than half the typical densities.**

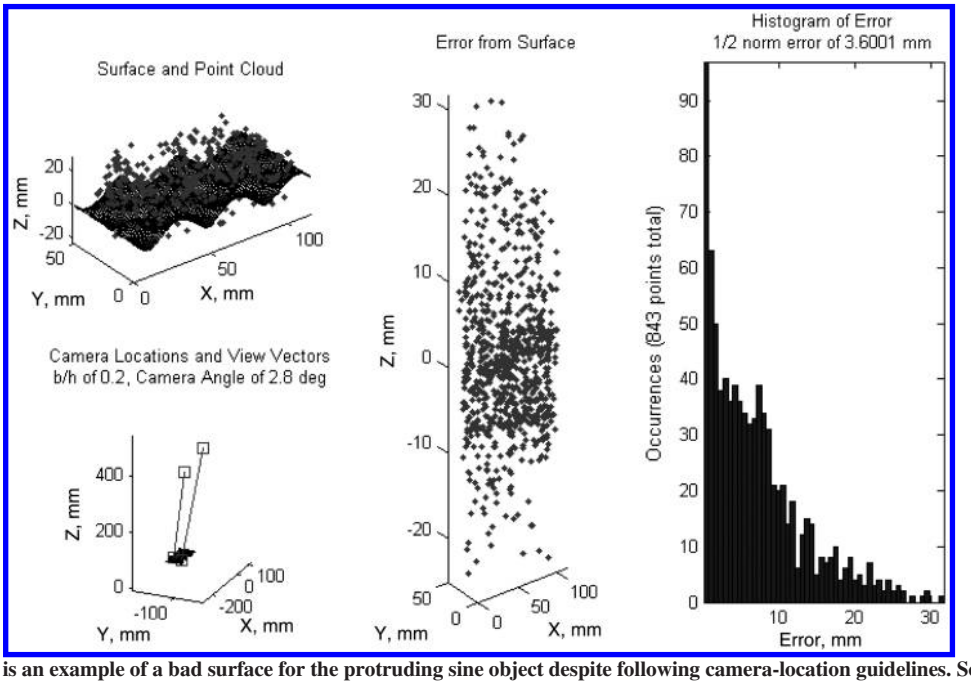


**Fig. 12 Shown here is an example of a bad surface for the protruding sine object despite following camera-location guidelines. So many of the points have a large error, and the true surface is unidentifiable by eye. Results could be refined via statistical methods, shown by the low 1/2 norm error, although they would need to be more sophisticated than simple outlier detection. On the left, results are shown over the truth surface with the relative camera locations shown in the lower plot. The center plot shows the error of the result by location, whereas the right-hand plot is a histogram of error occurrence.**

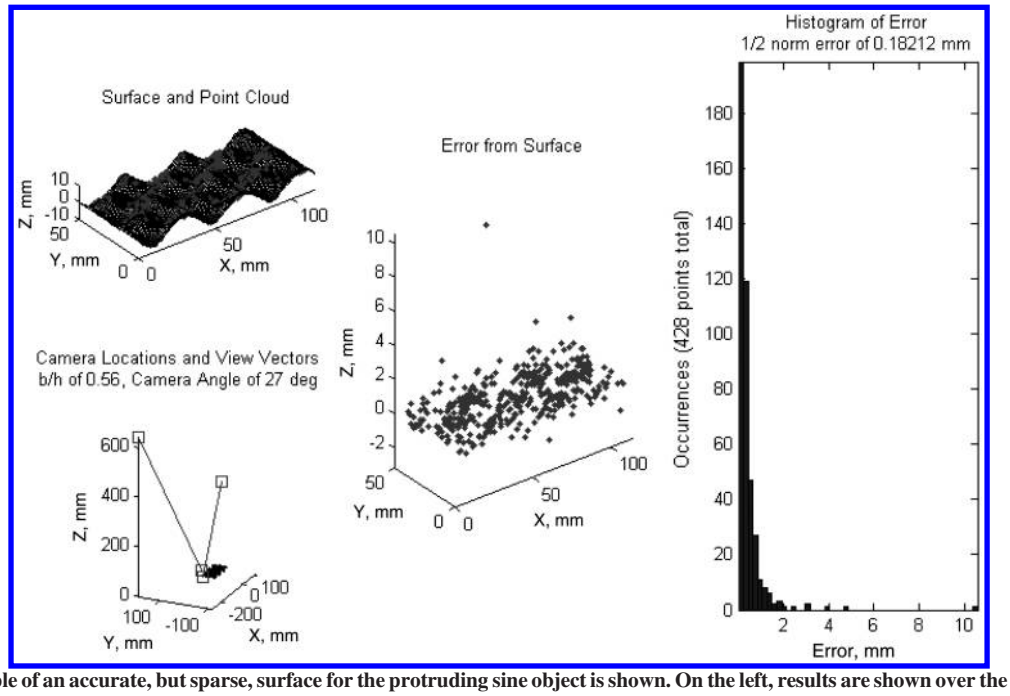


**Fig. 13 An example of an accurate, but sparse, surface for the protruding sine object is shown. On the left, results are shown over the truth surface with the relative camera locations shown in the lower plot. The center plot shows the error of the result by location, whereas the right-hand plot is a histogram of error occurrence. Texture photogrammetry guidelines are near the upper limit.**

then 2 deg, 1 mm, or 2% scaling. Because edges were more likely to have outliers, a 2.5 mm border was removed from the point cloud before registration so that these errant points would not affect the optimization.

Good meshes were typical, whereas some of the failures are shown to present how performance degrades. A graph of the percentiles shows the distribution of the median error of the meshes for each object in Fig. 8. For the gentle sine and planar surface, performance was normally better than 1:2000 (meaning the baseline distance was 2000 times larger than the error), whereas the protruding sine was closer to 1:1000. Note that the images were 2560 by 1920 pixels with the surface roughly having a length of only 900 pixels because of the need to view the coded targets in the background. Results show that accuracy can reasonably be estimated at pixel resolution. Despite thorough examination, no correlation between setup and accuracy could be found. Tables 2 and 3 condense results to percentiles for ranges of base-to-height ratios and camera angles. Each block lists the number of meshes and the accuracy (baseline per error) at three percentiles to show the spread of the accuracy between meshes. Results from the angular planes and the gentle sine were combined because they had similar performance. The suggested guidelines (camera angle $< 30$ deg and $0.1 \leq$ base-to-height $\leq 0.5$ [37]) did not show significantly better results, and there is slight evidence to suggest that a larger base-to-height ratio is better. The only correlations to results was that small camera angles or base-to-height ratios would produce more points and, therefore, denser meshes, as shown in Fig. 7. This section concludes with descriptions and examples of point clouds grouped by test object.

The gentle sine, Fig. 3a, gave accurate, dense point clouds without features deserving of comment. The planar test object, Fig. 3b, had one significant peak that noticeably affected the results. An example of a good mesh when not following texture photogrammetry guidelines is shown in Fig. 9. An example of a mesh on the tail-end of accuracy, but following texture photogrammetry guidelines, is shown in Fig. 10. Mean absolute surface of error by location did not show localization of the error for either object. Though error was evenly distributed, the number of points generated at a location was affected by features, including the angle of the surface for the planar case and the edge of the surface for both, as shown in Fig. 11. This is likely due to flat regions becoming entirely occluded, rather than the curved surfaces becoming gradually occluded.

The protruding sine test object, Fig. 3c, had noticeably worse performance, both in number of points produced and accuracy. With a greater chance of occlusions between any two images, fewer points

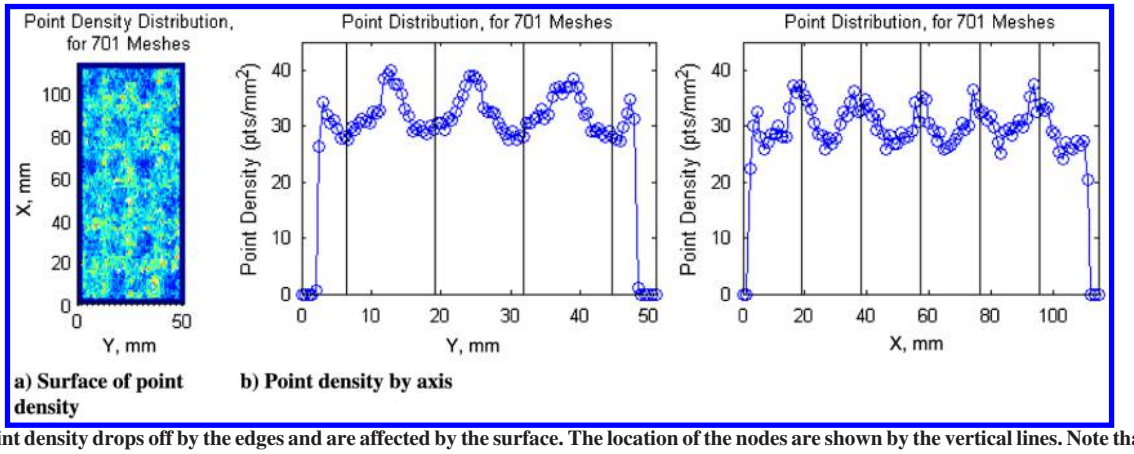


**Fig. 14 Point density drops off by the edges and are affected by the surface. The location of the nodes are shown by the vertical lines. Note that there is a higher point density at nodes in the low frequency direction ($X$) but lower density for the high frequency direction ($Y$).**

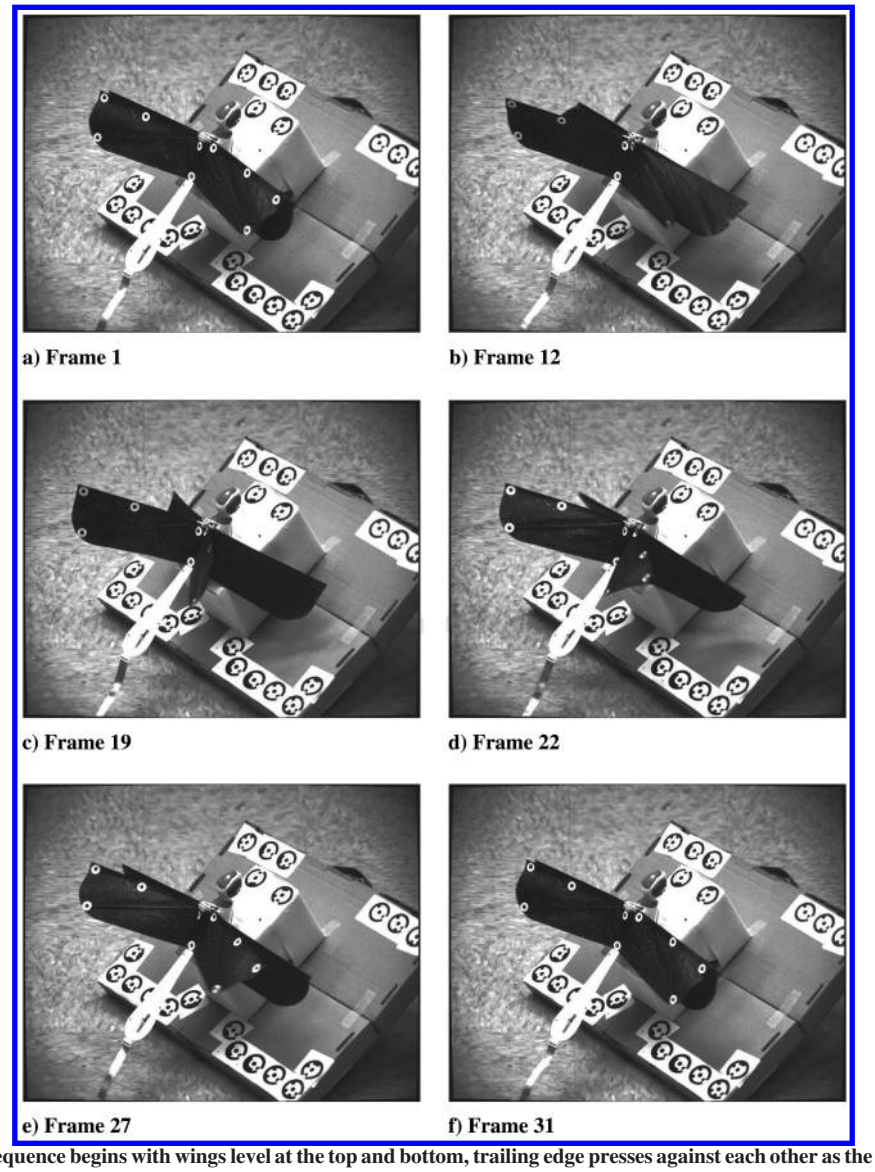


**Fig. 15 a) The flapping sequence begins with wings level at the top and bottom, trailing edge presses against each other as the wings go up as in 15b. The wing maintains this angle until the leading edge peaks at 15c. Then, the diagonal wing spar continues up at 15d, reversing the angle on the downward stroke as in 15e. 15f shows the return to the beginning of the sequence. Contrast has been adjusted for clarity of printouts.**

would be expected but not the doubling of point cloud error. Good meshes had noticeably more outliers, whereas bad meshes had no apparent surface, as shown in Fig. 12. Despite the apparent lack of a surface, the $1/2$ norm error is a quarter of the feature height, so advanced outlier detection methods could filter results to obtain a reasonable estimate of the surface. A new failure became noticeable on the protruding sine, that of sparse meshes in which not many points were generated. These meshes could have high or low accuracy. A mesh with only 428 points but low error is shown in Fig. 13. Error was again evenly distributed, and features affected the number of points generated at a location, as shown by the point density across all meshes of the protruding sine in Fig. 14.

In summary, results show accuracy of texture-based photogrammetry is similar to accuracy of target-based photogrammetry. Occlusion will cause fewer points to be created and can have an effect on accuracy. As with many dense point cloud methods, some outliers are typical, but the majority of points and meshes have errors on the order of 1:2000 or 1:1000. These results hold regardless of whether or not the suggested guidelines are followed.

## VI. Application

The results of Section V represent a near-ideal layout with many cameras, high-contrast speckle, a specific calibration, and a known surface. To show that texture-based photogrammetry can be applied with good results on natural systems, a high-speed capture was taken of a remote-controlled dragonfly in a stationary position flapping at 7.5 Hz. Two flexible legs were taped to a platform, allowing the dragonfly to bounce up and down and twist. Four 1.2 megapixel monochromatic cameras were synchronized at 225 fps and placed about the dragonfly. Zoom and focus are computer-controlled, and the calibration and uncertainty is addressed by Magree [40]. Figure 15 shows the motion of the wing through idealized photographs, whereas Fig. 16 shows the reconstructed 3-D wing shapes. An example of the images used for texture-based photogrammetry is shown in Fig. 17. The wing was nominally sampled every millimeter, and approximately 6000 points were generated per wing. The wing best matches the gentle sine shape from Section V. Taking a conservative accuracy estimate of 1:2000 and an image diagonal of 550 mm, the median error is expected to be less than 0.28 mm. This dense surface allows for detection of local surface properties, such as curvature. Distinct features, such as the edge of the region where the wings are pressed together, can be found with significantly higher resolution than previous target-based methods. Qualitative segmentation of the flap cycle is intuitive from being able to change the view of the 3-D wing shape. The results in Fig. 16 show the rising of the wing, the peak of the leading spar, and the maximum arch of the membrane along with the stages of the wings pressing together. The effect of the diagonal spar adding inertia and stiffness to the membrane is also seen. Mode shapes can be measured to arbitrary

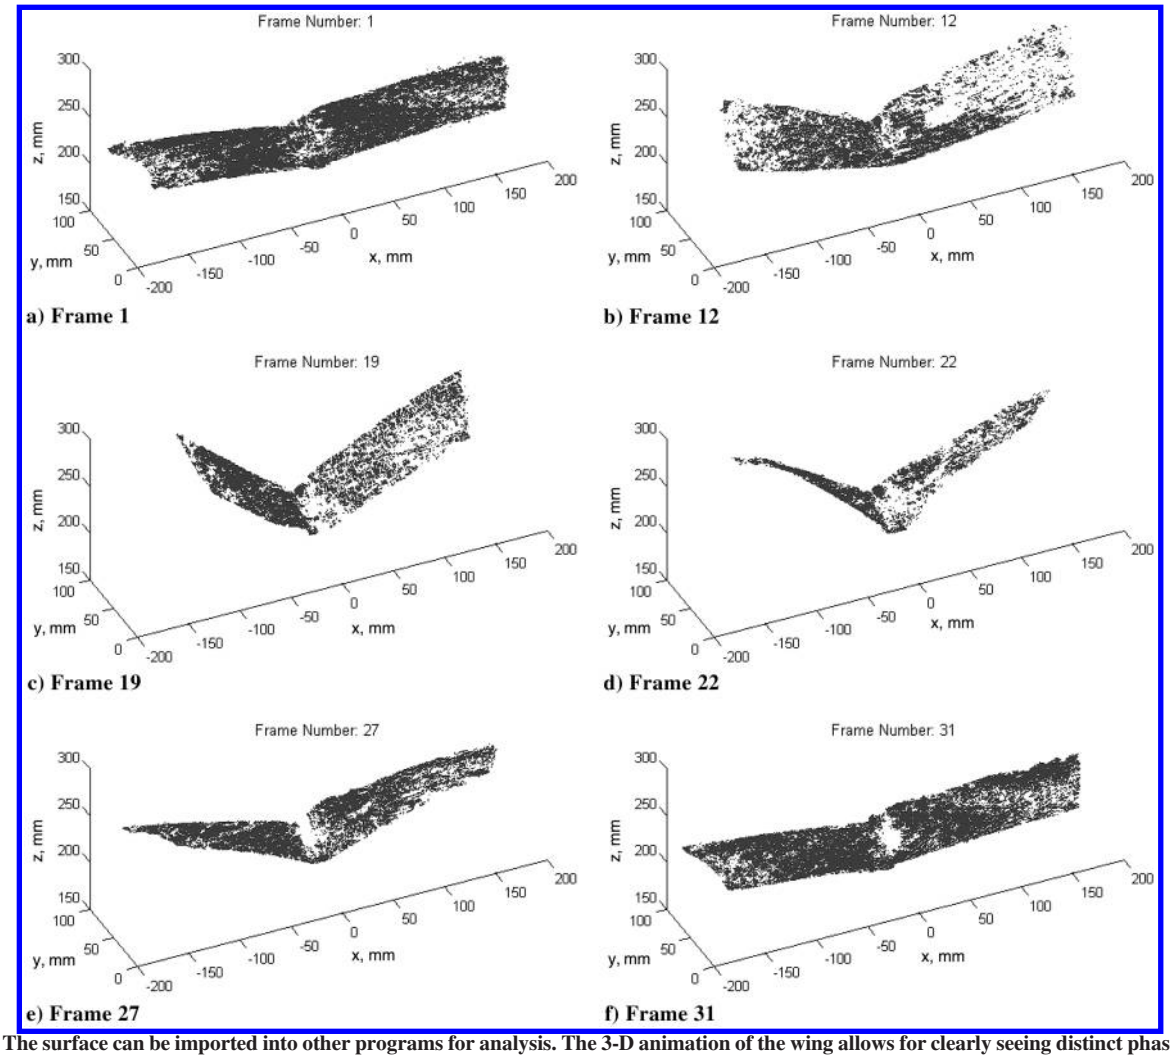


**Fig. 16 The surface can be imported into other programs for analysis. The 3-D animation of the wing allows for clearly seeing distinct phases listed in Fig. 15 because the view can be arbitrarily rotated.**

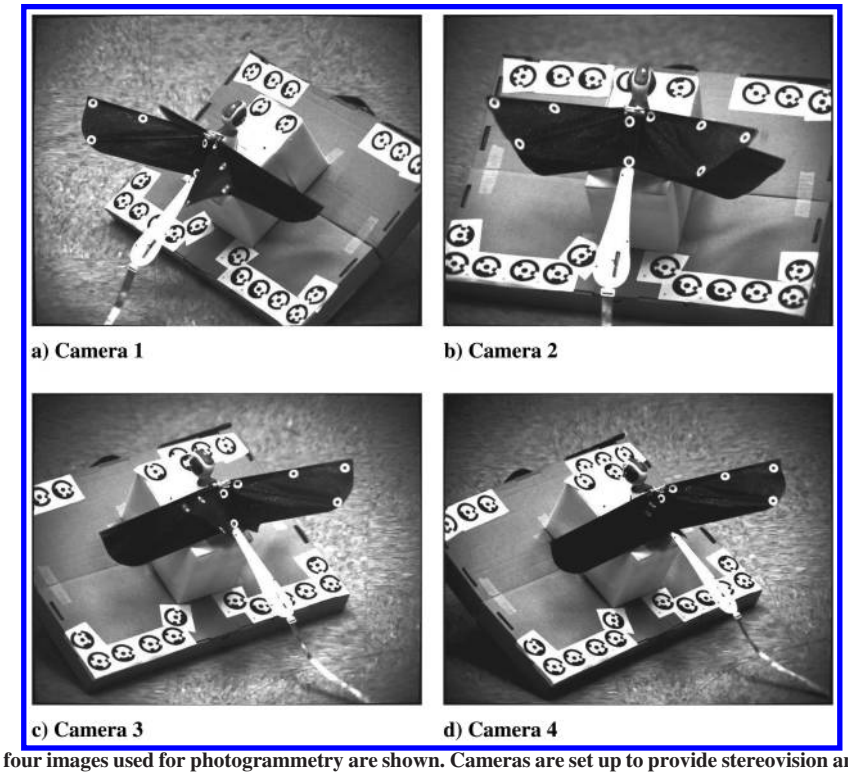


**Fig. 17 For frame 22, the four images used for photogrammetry are shown. Cameras are set up to provide stereovision and wide angles for accurately locating camera stations.**

resolution. High-density surface profiles, such as these, can be used to validate CFD analysis of deformable wings.

## VII. Conclusions

Texture-based photogrammetry stands as a significant improvement to the state of the art. Resolution of the surface is orders of magnitude greater than possible with target-based photogrammetry. By testing known profiles, practical accuracy has empirically determined error to be typically on the order of a pixel. The use of rounded surfaces shows that results are not artificially flattened, and the method does capture rolling details. The method of analysis is explained with the average error shown to be normally less than 1:2000 for surfaces without prominent protrusions and 1:1000 otherwise, which is on the order of the image pixel width. This accuracy compares well to what has been seen for target-based photogrammetry. Having validated accuracy on control surfaces, confidence in experimental results will be possible. Surprisingly, results show that texture-based photogrammetry is robust even when best practices regarding camera stations are not followed. The most significant correlation is that narrow camera views, or low base-to-height ratios, will generate more points on the surface.

High-speed video provides the wide view necessary to reconstruct the surface profile, has high-resolution for accurate dense surface profiles, and is fast enough to observe vibrations. Practical results show the ability to capture wing motion without blur and reconstruct the wing surface to a high resolution. This system will provide the tools needed to validate flapping wing analysis on relevant problems, such as maneuvering. Qualitative results derived from wing shapes can also assist in ad hoc ornithopter design.

## Acknowledgments

We would like to thank the Air Force Research Laboratory and the Air Force Office of Scientific Research for funding support, Ronald Tuttle and the Video Analysis and Content Extraction Lab for their help, and Joseph Blandino from the Virginia Military Institute for collaboration. Preliminary work by Daniel Magree, Gregory Briggs, and Chris Allen helped develop the tools and programs used in this work.

## References

[1] Qin, Z., and Librescu, L., "Static and Dynamic Validations of a Refined Thin-Walled Composite Beam Model," *AIAA Journal*, Vol. 39, No. 12, Aug. 2001, pp. 2422–2424. doi:10.2514/2.1259

[2] Wood, R., "The First Takeoff of a Biologically Inspired At-Scale Robotic Insect," *IEEE Transactions on Robotics*, Vol. 24, No. 2, Apr. 2008, pp. 341–347. doi:10.1109/TRO.2008.916997

[3] How, J. P., Bethke, B., Frank, A., Dale, D., and Vian, J., "Real-Time Indoor Autonomous Vehicle Test Environment," *IEEE Control Systems Magazine*, Apr. 2008, pp. 51–64. doi:10.1109/MCS.2007.914691

[4] Saad, E., Vian, J., Clark, G. J., and Bieniawski, S., "Indoor Multi-Vehicle Flight Testbed for Fault Detection, Isolation, and Recovery," *Infotech Aerospace Conference*, AIAA Paper 2009-1824, Seattle, WA, Apr. 2009.

[5] Ol, M., Parker, G., Abate, G., and Evers, J., "Flight Controls and Performance Challenges for MAVs in Complex Environments," *Guidance, Navigation, and Control Conference*, AIAA Paper 2008-6508, Honolulu, HI, Aug. 2008.

[6] Grzywna, J., Jain, A., Plew, J., and Nechyba, M., "Rapid Development of Vision-Based Control for MAVs through a Virtual Flight Testbed," *Proceedings of the 2005 IEEE International Conference*, IEEE Publications, Piscataway, NJ, Apr. 2005, pp. 3696–3702.

[7] Lienard, S., Johnston, J. D., Ross, B., and Smith, J., "Dynamic Testing of a Subscale Sunshield for the Next Generation Space Telescope (NGST)," *42nd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference and Exhibit*, AIAA Paper 2001-1268, Seattle, WA, Apr. 2001.

[8] Dharamsi, U. K., Evanchik, D. M., and Blandino, J. R., "Comparing Photogrammetry with a Conventional Displacement Measurement Technique on a 0.5 m Square Kapton Membrane," *43rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, AIAA Paper 2002-1258, Denver, CO, Apr. 2002.

[9] Leifer, J., Black, J. T., Smith, S. W., Ma, N., and Lump, J. K., "Measurement of In-Plane Motion of Thin-Film Structures Using Videogrammetry," *Journal of Spacecraft and Rockets*, Vol. 44, No. 6, Nov–Dec. 2007, pp. 1317–1325. doi:10.2514/1.25566

[10] Stahl, H. P., Alongi, C., Arneson, A., Bernier, R., Brown, B., Chaney, D., Cole, G., Daniel, J., Dettmann, L., Eng, R., Gallagher, B., Gareld, R., Hadaway, J., Johnson, P., Lee, A., Leviton, D., Magruder, A., Messerly, M., Patel, A., Reardon, P., Schwenker, J., Seilonen, M., Smith, K., and Smith, W. S., "Survey of Interferometric Techniques used to Test JWST Optical Components," *Proceedings of SPIE: The International Society for Optical Engineering*, Vol. 7790, 2010, pp. 779–002.

[11] Ma, C.-C., and Hsieh, D.-M., "Full-Field Experimental Investigations on Resonant Vibration of Cracked Rectangular Cantilever Plates," *AIAA Journal*, Vol. 39, No. 12, 2001, pp. 2419–2422. doi:10.2514/2.1258

[12] Wu, G. H., and Zeng, L. J., "Measuring the Kinematics of a Free-Flying Hawk-Moth (Macroglossum Stellatarum) by a Comb-Fringe Projection Method," *Acta Mechanica Sinica*, Vol. 26, No. 1, 2009, pp. 67–71. doi:10.1007/s10409-009-0306-y

[13] Markley, J., Stutzman, J., and Harris, E., "Hybridization of Photogrammetry and Laser Scanning Technology for As-Built 3-D CAD Models," *Proceedings of the IEEE Aerospace Conference*, IEEE Publications, Piscataway, NJ, Mar. 2008, pp. 1–10.

[14] Jennings, A. L., Black, J. T., Allen, C., Simpkins, J., and Sollars, R., "Vibrometer Steering System for Dynamic In-flight Tracking and Measurement," *Experimental Mechanics*, Vol. 51, 2010, pp. 71–84. doi:10.1007/s11340-010-9337-3

[15] Magree, D., Black, J., Jennings, A., Briggs, G., and Allen, C., "Pan-Tilt-Zoom Hybrid Camera System For Dynamic Tracking and Measurement," *AIAA Journal*, Vol. 49, No. 9, 2011, pp. 1988–1999, doi:10.2514/1.55040.

[16] Gruen, A., "Fundamentals of Videogrammetry: A Review," *Human Movement Science*, Vol. 16, 1997, pp. 155–187. doi:10.1016/S0167-9457(96)00048-6

[17] Yi, M., Liu, Z., and Peng, S., "Subpixel Edge Location Using Improved LSR," *Proceedings of IEEE Networking, Sensing and Control Conference*, IEEE Publications, Piscataway, NJ, Mar. 2005, pp. 970–975.

[18] Chang, C.-I., Chiang, S.-S., Du, Q., Ren, H., and Ifarragaerri, A., "An ROC Analysis for Subpixel detection," *IEEE International Geoscience and Remote Sensing Symposium*, Vol. 5, IEEE Publications, Piscataway, NJ, 2001, pp. 2355–2357.

[19] Sandwith, S., "Scale Artifact Length Dependence of Videogrammetry System Uncertainty," *Proceedings of SPIE: The International Society for Optical Engineering*, Vol. 3204, No. 36, 1997, pp. 36–47.

[20] Sandwith, S., "Real-Time Videogrammetry Uncertainty Computation and Test Results," *Proceedings of SPIE Conference on Three-Dimensional Imaging, Optical Metrology, and Inspection IV*, Vol. 3520, No. 111, SPIE, Bellingham, WA, Nov. 1998, pp. 111–112.

[21] Reid, D., and Lithopoulos, E., "High Precision Pointing System for Airborne Sensors," *IEEE Position Location and Navigation Symposium*, IEEE Publications, Piscataway, NJ, Apr. 1998, pp. 303–308.

[22] Lin, S.-Y., and Mills, J. P., "Integration of Videogrammetry and Terrestrial Laser Scanning for Dynamic Surface Modeling," *Proceedings of Videometrics VIII*, Vol. 5665, SPIE, Bellingham, WA, 2005, pp. 243–251.

[23] Black, J. T., Pitcher, N. A., Reeder, M. F., and Maple, R. C., "Videogrammetry Measurements of a Lightweight Flexible Wing in a Wind Tunnel," *Journal of Aircraft*, Vol. 47, No. 1, 2010, doi:10.2514/1.44545.

[24] Pappa, R. S., Black, J. T., Blandino, J. R., Jones, T. W., Danehy, P. M., and Dorrington, A. A., "Dot Projection Photogrammetry and Videogrammetry of Gossamer Space Structures," *Journal of Spacecraft and Rockets*, Vol. 40, No. 6, 2003, pp. 858–867. doi:10.2514/2.7047

[25] Harmon, R., Grauer, J., Hubbard, J., Conroy, J., Humbert, S., Sitaraman, J., and Roget, B., "Experimental Determination of Ornithopter Membrane Wing Shapes Used For Simple Aerodynamic Modelling," *26th AIAA Applied Aerodynamics Conference*, AIAA Paper 2008-6237, Honolulu, HI, Aug. 2008.

[26] Roget, B., Sitaraman, J., Harmon, R., Grauer, J., Conroy, J., Hubbard, J., and Humbert, S., "A Computational Study of Flexible Wing Ornithopter Flight," *26th AIAA Applied Aerodynamics Conference*, AIAA Paper 2008-6397, Honolulu, HI, Aug. 2008.

[27] Pitcher, N. A., and Maple, R. C., "A Static Aeroelastic Analysis of a Flexible Wing Mini Unmanned Aerial Vehicle," *38th Fluid Dynamics Conference*, AIAA Paper 2008-4057, Seattle, WA, 2008.

[28] Black, J. T., "New Ultra-Lightweight Stiff Panels for Space Apertures," Ph.d. Dissertation, Dept. of Mechanical Engineering, Univ. of Kentucky, Lexington, KY, 2006.

[29] Black, J. T., Smith, S. W., Leifer, J., and Bradford, L. J., "Experimental Characterization and Modeling of Dynamic Behavior of Semi-Rigid Thin Film Polyimide Panels," *Journal of Guidance, Control, and Dynamics*, Vol. 31, No. 3, 2008, pp. 490–500. doi:10.2514/1.32236

[30] Johnston, J. D., Blandino, J. R., and McEvoy, K. C., "Analytical and Experimental Characterization of Gravity Induced Deformation of Subscale Gossamer Structures," *Journal of Spacecraft and Rockets*, Vol. 43, No. 4, 2006, pp. 762–770. doi:10.2514/1.14298

[31] Pappa, R. S., Geirsch, L. R., and Quagliaroli, J. M., "Photogrammetry of a 5-m Inatable Space Antenna with Consumer-Grade Digital Cameras," *Experimental Techniques*, 2001, pp. 21–29. doi:10.1111/j.1747-1567.2001.tb00028.x

[32] Pappa, R. S., Lassiter, J. O., and Ross, B. P., "Structural Dynamics Experimental Activities in Ultra-Lightweight and in Atable Space Structures," *42nd Structures, Structural Dynamics, and Materials Conference*, AIAA 2001-1263, Seattle, WA, 2001.

[33] Leifer, J., Belvin, W., and Behun, V., "Evaluation of Shear Compliant Borders for Wrinkle Reduction in Thin Film Membrane Structures," *44th Structures, Structural Dynamics, and Materials Conference*, AIAA Paper 2003-1984, Norfolk, VA, 2003.

[34] Black, J. T., Leifer, J., DeMoss, J. A., and Walker, E. N., "Experimental and Numerical Correlation of Gravity Sag in Solar Sail Quality Membranes," *Journal of Spacecraft and Rockets*, Vol. 44, No. 3, 2007, pp. 522–527. doi:10.2514/1.20958

[35] Blandino, J. R., Pappa, R. S., and Black, J. T., "Modal Identication of Membrane Structures with Videogrammetry and Laser Vibrometry," *44th Structures, Structural Dynamics, and Materials Conference*, AIAA Paper 2003-1745, Norfolk, VA, 2003.

[36] Bellmann, A., Hellwich, O., Rodehorst, V., and Yilmaz, U., "A Benchmarking Dataset for Performance Evaluation of Automatic Surface Reconstruction Algorithms," *Proceedings of the IEEE Computer Vision and Pattern Recognition*, IEEE Publications, Piscataway, NJ, June 2007, pp. 1–8.

[37] *PhotoModeler Pro v6 User Manual*, Eos Systems, Inc., Vancouver, BC, Canada, 2008.

[38] Gonzalez, R. C., and Woods, R. E., *Digital Image Processing*, 3rd ed., Prentice–Hall, Upper Saddle River, NJ, 2008.

[39] Olague, G., and Mohr, R., "Optimal Camera Placement for Accurate Reconstruction," *Pattern Recognition*, Vol. 35, No. 4, 2002, pp. 927–944. doi:10.1016/S0031-3203(01)00076-0

[40] Magree, D. P., *A Photogrammetry-Based Hybrid System for Dynamic Tracking and Measurement*, Master's, Thesis, Air Force Inst. of Technology, Wright-Patterson Air Force Base, OH, June 2010.

N. Wereley
*Associate Editor*

**This article has been cited by:**


1. Philip Sapirstein. 2018. A high-precision photogrammetric recording system for small artifacts. *Journal of Cultural Heritage* **31**, 33-45. [Crossref]
2. P. Sapirstein. 2016. Accurate measurement with photogrammetry at large sites. *Journal of Archaeological Science* **66**, 137-145. [Crossref]
3. Xiaoliang Gong, Stephan Bansmer. 2015. Horn–Schunck optical flow applied to deformation measurement of a birdlike airfoil. *Chinese Journal of Aeronautics* **28**:5, 1305-1315. [Crossref]
4. Takashi IWASA, Naoko KISHIMOTO, Ken HIGUCHI, Motoharu FUJIGAKI, Nozomu KOGISO. 2015. Proposal of Precise Shape Measurement Method Synthesizing Surface Measurement and Point Measurement for Large Space Structures. *AEROSPACE TECHNOLOGY JAPAN, THE JAPAN SOCIETY FOR AERONAUTICAL AND SPACE SCIENCES* **14**:0, 95-103. [Crossref]
5. Daniel D. Doyle, Alan L. Jennings, Jonathan T. Black. 2014. Optical flow background estimation for real-time pan/tilt camera object tracking. *Measurement* **48**, 195-207. [Crossref]
6. Daniel D. Doyle, Alan L. Jennings, Jonathan T. Black. Optical flow background subtraction for real-time PTZ camera object tracking 866-871. [Crossref]