



Full-Scale Supersonic Parachute Shape Reconstruction Using Three-Dimensional Stereo Imagery

Jason Rabinovitch,* Gregory S. Griffin,† William Seto,‡ Clara O'Farrell,§ Christopher L. Tanner,|| and Ian G. Clark**

Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California 91109

<https://doi.org/10.2514/1.A34717>

This paper describes three-dimensional (3D) stereo reconstruction efforts for the full-scale supersonic parachute deployment tests performed under the NASA Advanced Supersonic Parachute Inflation Research Experiment (ASPIRE) program. Specific time-resolved data products that have been generated for the ASPIRE tests and are presented in this work include parachute projected area, parachute vent-hole location, and parachute surface reconstruction (3D point clouds).

I. Introduction

THE supersonic Disk-Gap-Band (DGB) parachute was developed during a series of campaigns undertaken by NASA in the 1960s and 1970s for the Viking project [1] that included the Planetary Entry Parachute Program (PEPP), Supersonic Planetary Entry Decelerator (SPED), Supersonic High Altitude Parachute Experiment (SHAPE), and Balloon Launched Decelerator Test (BLDT) supersonic flight test campaigns [2–4]. Following the successful landing of the Viking spacecraft in 1976, every U.S. mission to the Martian surface has used a variant of the Viking DGB: 1997s Mars Pathfinder [5,6], the Mars Polar Lander two years later [7], the twin Mars Exploration Rovers in 2004 [8], the Phoenix lander in 2007 [9], 2012s Mars Science Laboratory (MSL) [10], and the InSight lander in November 2018 [11]. Recent experience with supersonic parachute testing has suggested that the strength margins in more recent DGB designs are not as robust as previously thought [12], underscoring the need for further testing of supersonic parachutes and improved understanding of the stresses on an inflating canopy.

The Advanced Supersonic Parachute Inflation Research Experiments (ASPIRE) project was initiated in 2016 to address these needs. An initial sequence of sounding rocket flight tests took place between October 2017 and September 2018, and resulted in the successful deployment of 21.5 m DGBs at Mach numbers between 1.77 and 1.97 at an altitude of approximately 40 km [13,14]. The test platform included three cameras for stereographic reconstruction of the canopy shape during inflation. This paper will describe efforts to reconstruct the three-dimensional (3D) canopy shape during inflation based on stereo image data acquired by the onboard cameras.

Video recordings captured by a single camera mounted on the test vehicle and trained on the parachute have been a part of parachute development tests from the early PEPP [2,15], SPED [3,16], SHAPE

[4,17,18], and BLDT tests [19–21] to the recent Low-Density Supersonic Decelerators (LDSD) [22] and Orion Crew Exploration Vehicle (CEV) Parachute Assembly System (CPAS) [23] ringsail test campaigns.†† In addition to providing situational awareness and qualitative information about the parachute deployment process, these recordings were used to estimate the projected area of the canopy during inflation and the position of the canopy relative to the test vehicle. However, because a single camera was used, these estimates relied on assumptions of the canopy trailing distance and could not provide an inflated canopy shape. The reliance on uplook cameras posed additional challenges for photogrammetric analysis, as the inflating canopy was observed through the suspension lines and other parachute rigging. In addition, portions of the image could be occluded by other elements of the vehicle or parachute assemblies (thermal blanketing, parachute mortar lid, etc.).

As a result, experimental studies of canopy geometry, fabric stress, or fluid–structure interaction have relied on imagery recorded from (often multiple) side- or downstream-view angles. In recent Orion CPAS tests, images from a side-view camera mounted on a chase helicopter were used to supplement the onboard camera imagery in reconstructing the inflated shape and volume of the main parachutes [24]. Brent and DeWeese [25] also used side-view imagery obtained during low-altitude deployments of flat-circular canopies to develop models for the time-varying shape of inflated canopies of this type. Houmar [26] used these models, in conjunction with projected area measurements from the Viking DGB development tests to model fabric stress on the Viking DGB during inflation.

The reliance on side or downstream camera views to reconstruct canopy geometry also extends to wind tunnel testing. While Schoenenberger et al. [27], Underwood et al. [28], and Cruz et al. [29] tracked the motion of subscale parachute canopies in a wind tunnel using a single camera located directly downstream of the canopies, Cruz et al. used additional images captured using a side-view camera to reconstruct the average inflated canopy shape. Mayhue [30] used four cameras mounted on the walls of the Arnold Engineering Development Complex (AEDC) Propulsion Wind Tunnel to track the motion of subscale DGB canopies and reconstruct their projected area. In several subsonic tests of full-scale (or near-full-scale) parachutes at the National Full Scale Aerodynamics Complex, an array of cameras positioned around the test section were used to track the parachutes' motion and reconstruct their projected area [31]. Of note, the tracking was aided by the use of photorefective targets attached to the canopy fabric. Similarly, photorefective targets were used to aid the measurement of the projected area during inflation of model DGBs deployed at Mach numbers ranging from 0.5 to 3.0 at the Glenn Research Center 10' × 10' tunnel [32].

In the ASPIRE test program, the high velocities and elevations at parachute deployment precluded the use of chase vehicles to obtain side-view images of the parachute during supersonic flight. In addition, the supersonic deployment environment and the desire to

Presented as Paper 2019-1629 at the AIAA SciTech 2019 Forum, San Diego, CA, January 7–11, 2019; received 24 November 2019; revision received 2 April 2020; accepted for publication 8 April 2020; published online 18 May 2020. Copyright © 2020 by the American Institute of Aeronautics and Astronautics, Inc. The U.S. Government has a royalty-free license to exercise all rights for Governmental purposes. All other rights are reserved by the copyright owner. All requests for copying and permission to reprint should be submitted to CCC at www.copyright.com; employ the eISSN 1533-6794 to initiate your request. See also AIAA Rights and Permissions www.aiaa.org/randp.

*Mechanical Engineer, EDL and Formulation Group, 4800 Oak Grove Drive, M/S: T1708-112. Senior Member AIAA.

†Robotics Technologist, Aerial and Orbital Image Analysis Group, 4800 Oak Grove Drive, M/S: 198-219.

‡Robotics Technologist, Maritime and Multi-Agent Autonomy Group, 4800 Oak Grove Drive, M/S: 198-219.

§Guidance and Control Engineer, EDL Guidance & Control Systems Group, 4800 Oak Grove Drive, M/S: 321-220. Member AIAA.

||Mechanical Engineer, EDL and Formulation Group, 4800 Oak Grove Drive, M/S: 321-220. Member AIAA.

**Systems Engineer, EDL Systems & Advanced Technologies Group, 4800 Oak Grove Drive, M/S: 321-400. Member AIAA.

††The LDSD project used a pair of cameras for redundancy and to double the effective frame rate. However, the camera fields of view were near identical.

test the parachutes in their flight configuration made the placement of fiducial markers on the canopy unfeasible. As a result, a system of three synchronized high-speed, high-resolution onboard uplook cameras was developed for ASPIRE. The baseline separation between the cameras was maximized within the constraints of the sounding rocket platform and a painstaking calibration and pre-processing method was developed to reduce measurement errors. Feature extraction and matching algorithms were used to reconstruct the 3D position of points on the canopy during inflation and supersonic flight.

This paper describes and shows results from the first steps taken to process this unique data set, and it is understood that future efforts will continue to produce high-fidelity data sets from the data presented in this work. This paper is organized as follows: Sec. II provides an overview of the ASPIRE test campaign, Sec. III describes the ASPIRE onboard instrumentation and data sets acquired during the flights, Sec. IV describes the current image processing techniques used, and Sec. V shows the preliminary data products calculated from the data sets, followed by a discussion and conclusions.

II. ASPIRE Test Campaign Overview

The ASPIRE project was started in 2016 to develop NASA's capability to test supersonic parachutes at Mars-relevant conditions using a sounding rocket test platform. The initial series of ASPIRE flights were focused on testing two candidate parachute designs for the Mars 2020 project: a build-to-print version of the Mars Science Laboratory (MSL) DGB parachute and a strengthened parachute with the same geometry but stronger materials [33]. The parachutes were delivered to targeted deployment conditions representative of flight at Mars by NASA's Sounding Rocket Operations Contract (NSROC) sounding rockets. Three flights took place between October 2017 and September 2018. During the first test (SR01) [13], a build-to-print DGB was tested at conditions similar to those that MSL recorded during Mars entry, descent, and landing (EDL). During the second and third tests (SR02 and SR03, respectively), the strengthened parachute was tested at incrementally higher loads [14]. All three flights were successful, resulting in minimal damage to the parachute test articles.

The concept of operations for the ASPIRE sounding rocket tests is outlined in Fig. 1. The sounding rocket assembly, consisting of a Terrier first stage, a Black Brant second stage, and the roughly 1200 kg payload section containing the experiment, is launched out of NASA's Wallops Flight Facility (WFF). The system is rail-launched and spin-stabilized at 4 Hz. During flight, the first and second stages burn out at

altitudes of approximately 700 m and 16.7 km, respectively, and the payload section reaches apogee between 50 and 55 km. When the payload reaches the target dynamic pressure q_∞ and Mach number conditions, the parachute is mortar-deployed. The deployment, inflation, and supersonic and subsonic aerodynamics of the parachute are analyzed by a suite of instruments including a three-camera high-speed/high-resolution stereographic video system trained on the parachute, situational awareness video cameras, a set of load pins at the interface of the parachute triple-bridle and the payload, and a GPS and inertial measurement unit (IMU) onboard the payload. After decelerating to subsonic speed, the parachute and payload descend to the ocean for recovery and inspection.

Figure 2 shows a schematic of the ASPIRE configuration prior to launch from WFF. The payload section has a maximum diameter of 0.72 m, and is 7.54 m in length. It comprises a nose cone ballast section that is jettisoned before splashdown, a section housing electronics and foam for buoyancy, a telemetry section, a section housing the attitude control system (ACS), and the experiment section containing the parachute mortar tube and ASPIRE instrumentation. The payload also includes an aft-transition section that remains attached to the Black Brant motor at separation. After separation from the Black Brant, the mass of the payload, including the stowed parachute and 600 kg nose cone, is approximately 1200 kg.

Figure 3 shows a schematic of the ASPIRE configuration after parachute deployment. The relevant dimensions of the parachute-payload system are labeled in the schematic, and their values are also listed in Table 1. The MSL build-to-print parachute tested on SR01 was an 80-gore DGB with a nominal diameter D_0 of 21.35 m. The majority of the canopy is constructed using a Nylon fabric with a rated strength of 74 N/cm (42 lbf/in.). The parachute is built using a continuous line construction where the suspension lines continue into the canopy as the radials, cross the vent, and continue along the opposite side of the canopy. The radials/suspension lines are constructed from 9.3 kN (2100 lbf) Technora line. The entire packed parachute assembly has a mass of 55.8 kg.

The strengthened parachute tested on SR02 and SR03 is an 80-gore DGB with a design nominal diameter D_0 of 21.45 m. The geometry of the strengthened canopy is identical to that of the build-to-print DGB, but higher-strength materials are used in its construction. The broadcloth gores are constructed using a custom Nylon fabric with a rated strength of 190 N/cm (110 lbf/in.). The radials/suspension lines are constructed from 14 kN (3200 lbf) Technora line, and the entire packed parachute assembly has a mass of 85 kg.

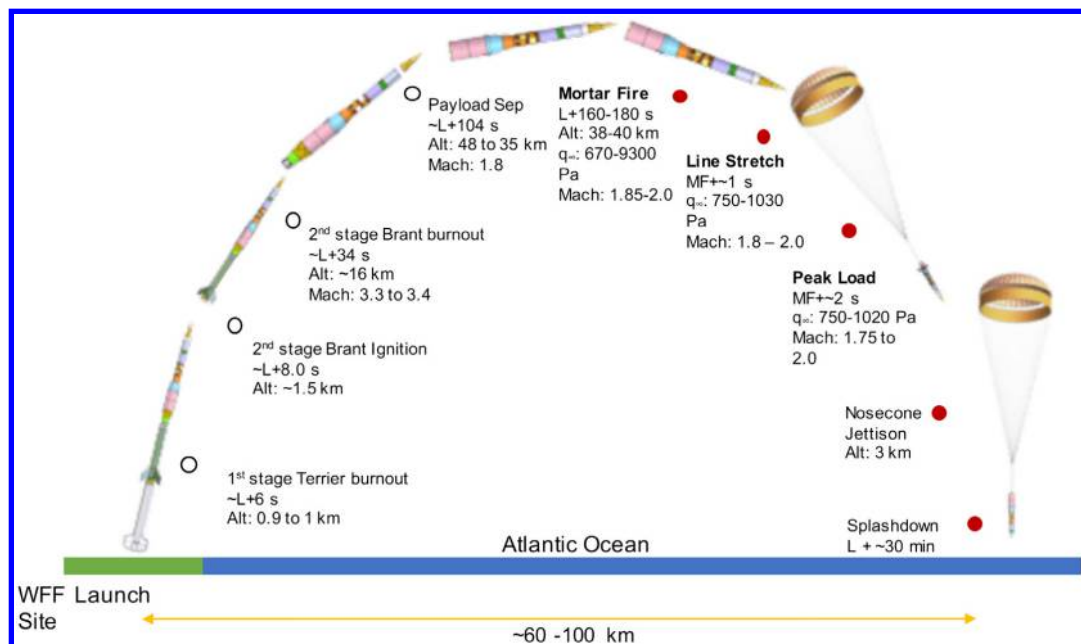


Fig. 1 ASPIRE concept of operations.

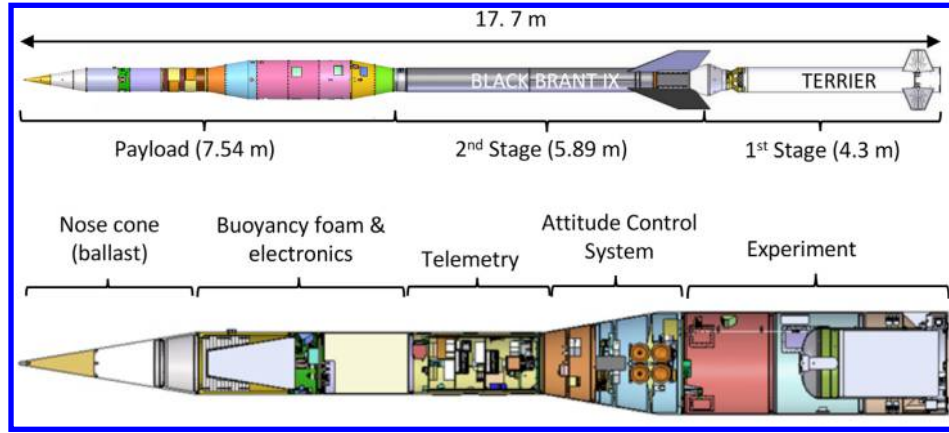


Fig. 2 Schematic of the ASPIRE launch configuration and payload.




Fig. 3 ASPIRE parachute system.

Table 2 shows the conditions at the beginning of parachute inflation for the three flights. Inflation is considered to start at suspension line stretch, i.e., when the entirety of the canopy exits the bag. Table 2 lists the wind-relative Mach number and dynamic pressure at line stretch. The time to full inflation is defined as the time from line stretch until the point when the parachute first achieves peak load. This first peak in parachute load may not be the highest peak, as supersonic parachutes often show a collapse and re-inflation. The last column in Table 2 shows the peak measured parachute load, regardless of when it occurred.

Table 1 As-designed dimensions of the ASPIRE parachutes

Item	Symbol	SR01	SR02 and SR03
Parachute reference diameter (m)	D_0	21.35	21.45
Parachute reference area (m ²)	S_0	358.00	361.36
Vent diameter (m)	D_V	1.50	1.50
Disk diameter (m)	D_D	15.44	15.62
Gap height (m)	H_G	0.90	0.90
Band height (m)	H_B	2.60	2.60
Geometric porosity (%)	λ_g	12.5	12.8
Suspension line length (m)	L_S	36.47	36.47
Riser length (m)	L_R	7.06	7.78
Bridle length (m)	L_B	1.35	1.30
Forebody diameter (m)	d	0.72	0.72
Packed mass (kg)	m	56.2	82.1

III. Experimental Data

The ASPIRE payload is equipped with scientific instrumentation to measure the trajectory, aerodynamics, and performance of the test vehicle and test article. A summary of the onboard instrumentation is provided in Table 3. The trajectory of the payload during each test is reconstructed using the measurements from the GLN-MAC IMU, the onboard GPS receiver, and three ground-based radars by means of an extended Kalman filter [34,35]. In addition, meteorological balloons carrying radiosondes are released hourly during launch operations. These provide measurements of the temperature, density, and winds up to 40 km.

Figure 4 shows the aft-most portion of the payload, which houses the parachute mortar. A specially designed instrumentation ring, also shown in Fig. 4, fits around and above the parachute mortar tube, and houses many of the key instruments. Three load pins are installed on the ring, at the locations where the parachute triple bridle legs attach to the payload, to measure the tension on the parachute riser. The load pins are rated to 400 kN (90 klbf), but calibrated to peak loads of 133 kN (30 klbf) for SR01 and SR02, or 178 kN (40 klbf) for SR03 to reduce the uncertainty in the load pin measurements. The total aerodynamic force on the parachute (which includes the measured tension and the portion that contributes to the deceleration of the parachute) was then derived from the load pin and IMU measurements.

Two sets of cameras are used to record the parachute deployment and inflation events: a set of high-resolution, high-speed (HS) cameras, and a set of situational awareness cameras. The HS camera system consists of three IDT OS-10 cameras paired with a ruggedized Schneider APO-Xenoplan 2.0/20 mm compact C-mount lens. The cameras are mounted around the circumference of the instrumentation ring (which has a radius of 29.6 cm at the camera mounting locations), at 120° intervals. They record 12-bit grayscale images at 1000 frames per second at a 4K resolution. Recording is performed with an aperture of 4.0 and a 40 μ s exposure time.

The situational awareness cameras are a set of standard GoPro Hero 4 cameras. While the footage from these cameras is not used for canopy shape reconstruction (the images are not shown in this work), the GoPros provides a wider-angle contextual view, an National Television System Committee (NTSC) feed that was tele-metered to ground during the flights for situational awareness, and a

Table 2 Parachute deployment conditions

Test	Mach at line stretch	q_∞ at line stretch, Pa	Density at line stretch, kg/m ³	Static pressure at line stretch, Pa	Inflation time, s	Peak load, kN
SR01	1.79	492	3.05×10^{-3}	218	0.506	144
SR02	2.00	745	3.68×10^{-3}	266	0.456	248
SR03	1.88	1028	6.02×10^{-3}	415	0.410	300

Table 3 Key instrumentation summary

Device	Sample rate/resolution	Notes
GLN-MAC IMU	400 Hz	Provides both 100 and 400 Hz raw data
Javad TR-G2 GPS	20 Hz	
Radars (×3)	50 Hz	
Load pins ^a	1 kHz	400 kN rated, calibrated to 133 or 178 kN ^b
HS cameras (×3) ^a	1000 fps, 3840 × 2400	
Situational video (×2) ^a	120 fps, 1920 × 1080	Global shutter, 12-bit grayscale images
Situational video (×1) ^a	30 fps, 3840 × 2160	Rolling shutter
Radiosondes (×4)	1 Hz	SR02 and SR03 only; rolling shutter
		Balloon-launched every hour

^aSee Fig. 4.^bCalibrated to 133 kN on SR01 and SR02, and to 178 kN on SR03.

longer recording of the parachute during descent. Two of these cameras were set to record at 1080p and 120 fps, and to telemeter video over NTSC during flight. On SR02 and SR03, a third GoPro was set to record 4K video at 30 fps.

As the ASPIRE test campaign was intended to certify a parachute for use on the Mars 2020 mission, the configuration of the test articles is identical to that of the flight articles intended for use at Mars. Fiducial markers and rigidized elements for distance calibration are not added to the test articles because their presence could affect stress distributions in the canopy and general parachute performance, and this could compromise the validity of the ASPIRE test results. Speckle patterns in the canopy broadcloth were also considered but ultimately not included because the effect of dyes or paints on the material properties (especially after undergoing the required environmental testing for a Mars mission) is unknown.

IV. Image Processing

This section provides an overview of the image processing techniques used to produce 3D point clouds representing the surface of the parachute canopy. These point clouds are generated from the high-speed images taken during the ASPIRE flights.



A. Overview

Figure 5 illustrates some of the difficulties encountered when using the three HS cameras for 3D reconstruction of the parachute surface during flight. The parachute surface, which lacks fiducials or rigid elements for distance calibration, twists and folds at angles that can be nearly perpendicular to the camera line of sight. Uncontrolled lighting conditions during flight lead to dark regions and specular reflections, which further complicate the analysis [36]. Small gain variations in individual pixels must be corrected to accurately establish correspondences between images.

Out of the various factors that limit surface coverage and resolution, several are found to be particularly significant:

1) Occlusions: Fast-moving suspension lines and other support structures obstruct the view of the parachute surface during the flights. These foreground structures can be particularly confusing for many popular dense stereo reconstruction methods that rely on template-based disparity searches over epipolar lines in rectified image pairs.

2) Complex topologies are created during deployment with surfaces that remain nearly perpendicular to the camera line of sight (especially near the canopy skirt), which create surfaces that are difficult to reconstruct accurately.

**Fig. 4** Instrumentation ring and sensor locations.**Fig. 5** Challenges for 3D reconstruction of the parachute surface include 1) occlusions, 2) complex topologies, 3) uncontrolled lighting conditions, and 4a) variations in pixel gain over the sensor and 4b) after flat-field gain correction.

3) Variations in pixel gain over the sensor require a flat-field gain correction algorithm to be added into the data-processing procedure.

4) Camera movement: Despite efforts to stabilize the camera mounts, it is believed that high g -forces lead to small but significant changes in camera pose that must be corrected during image post-processing.

5) Camera baselines: Three HS cameras equally spaced around the instrumentation ring form three baselines with a length of 0.514 m. The surface reconstruction error varies inversely with this baseline distance [37], which is constrained by the radius of the sounding rocket.

In an indoor wind-tunnel environment under controlled lighting conditions, camera spacing can be increased as needed along stable mounts located *behind* the parachute, where suspension lines and other obstructions are not visible [38]. For the purpose of the current in-flight surface reconstruction efforts, an analysis workflow is adopted that minimizes the sources of error that cannot be controlled.

As shown in Fig. 6, the workflow begins with a preflight flat-field calibration. By placing gray and white cards of uniform brightness in front of each camera lens, the resulting out-of-focus images are used to establish a map of gain fluctuations over individual sensor pixels. This correction significantly increases the number and quality of correspondences available for pose correct and point-cloud reconstruction, as described in Sec. IV.C. For both the preflight and in-flight imagery, GPS timestamps are used to resequence the frames to correct for a small (<1 per 3000) fraction of dropped images. The resulting image sets are synchronized to within 13, 25, and 16 μ s for SR01, SR02, and SR03, respectively.

B. Camera Calibration

Preflight camera calibration tests are performed near the launch facility. With the instrumentation ring pointed toward the horizon, a

flat target board with regularly spaced dots is carried across the overlapping field of view of the HS cameras. Care is taken to image these boards at a variety of locations and distances, covering as much of the focal plane of each sensor as possible. A common analysis procedure is used whereby dot positions are extracted in each image and triangulated across N sets of image triplets to produce 3D coordinates in space. These coordinates, in turn, are reprojected back onto the individual image planes for comparison with the original dot positions [39–42]. A least-squares fit procedure optimizes the camera intrinsic and extrinsic parameters to minimize the overall reprojection error in the image plane. Because each 10×10 array of dots is constrained to lie along a regular planar grid, the degrees of freedom in the fit are drastically reduced, from 300 N dot coordinates to 6 N coordinates representing the individual board positions and orientations.

The specific implementation in this work is performed in MATLAB and uses CAHVOR models [43,44] with eight intrinsic parameters per camera, and six additional extrinsic parameters specifying the positions and orientations of cameras 2 and 3 with respect to camera 1. Therefore a total of 36 model parameters are minimized using dots extracted from 80 to 120 images of target boards placed at distances between 3 and 60 m. The final mean reprojection error over all dots in all images varied from 0.29 to 0.34 pixels in the different calibrations performed. Table 4 gives the resulting field of view and the relative distance and orientation of the three cameras, where two intrinsic and two extrinsic parameters are listed: horizontal and vertical field of view (hFOV and vFOV, respectively), and relative distance and orientation (d and θ , respectively) relative to camera 1.

C. Pose Estimation and Point Cloud Reconstruction

This work implements a method that uses matched feature key points to correct camera extrinsics, while simultaneously reconstructing the parachute canopy. Figure 7 illustrates the process that starts

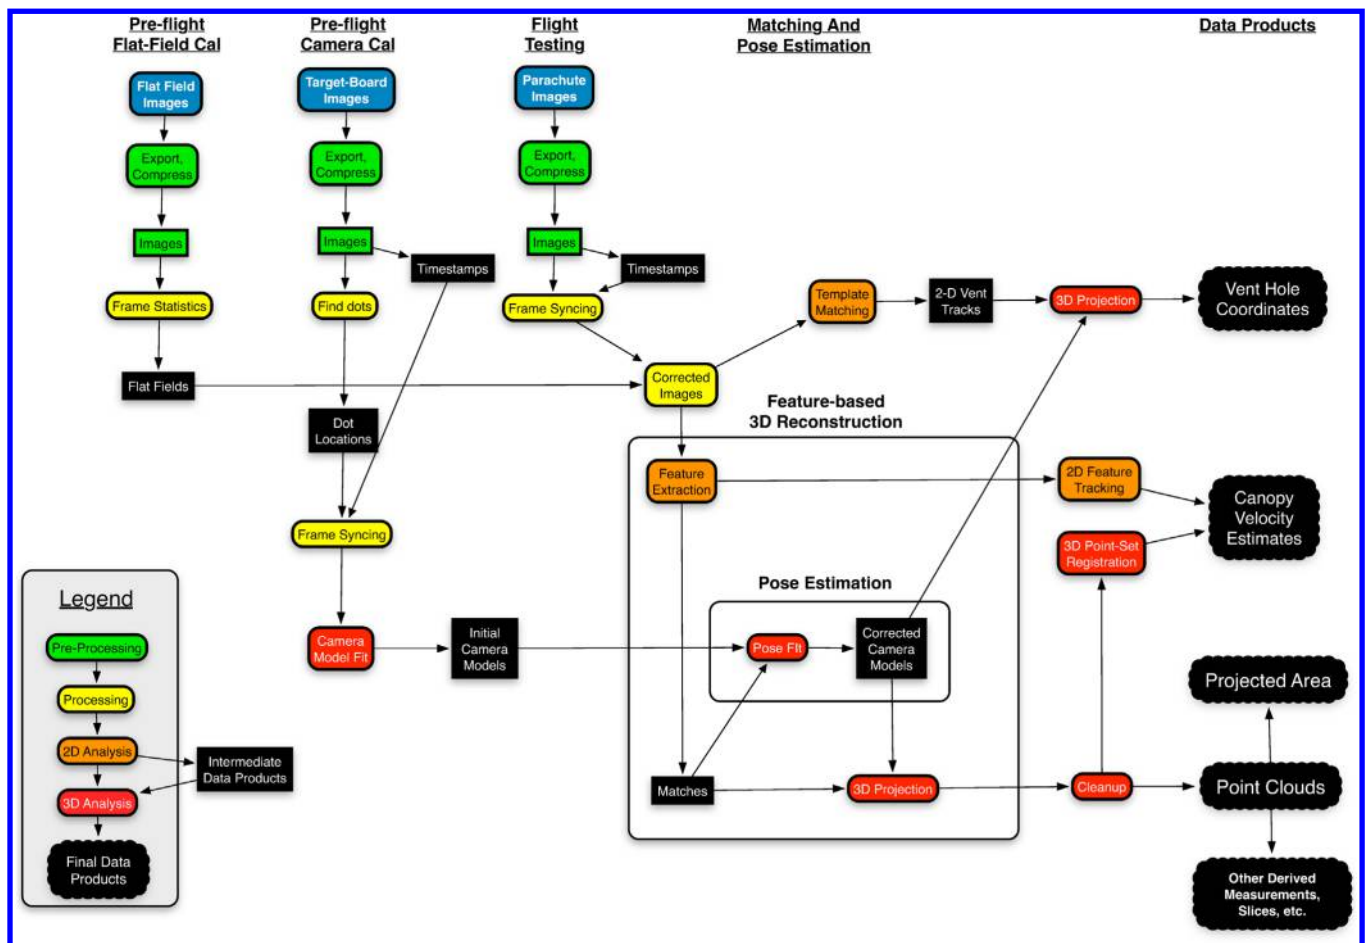


Fig. 6 Image analysis workflow.

Table 4 Results of preflight camera calibration for the three HS cameras used in SR01, SR02, and SR03

	HS camera 1		HS camera 2				HS camera 3			
	hFOV, °	vFOV, °	hFOV, °	vFOV, °	d_{12} , m	θ_{12} , °	hFOV, °	vFOV, °	d_{13} , m	θ_{13} , °
SR01	47.7	31.4	47.7	31.5	.514 m	60.1	47.7	31.4	.514 m	−60.0
SR02	47.9	31.6	47.8	31.6	.513 m	60.0	47.9	31.6	.513 m	−60.2
SR03	47.8	31.4	47.8	31.5	.514 m	60.2	47.9	31.5	.514 m	−60.1

with extracting and matching Scale Invariant Feature Transform (SIFT) key point and features [45] between image pairs using the MATLAB toolbox in VLFEAT [46]. The problem of using image correspondences to estimate pose while constructing detailed point clouds has been well studied [47]. The current software is adapted from routines developed by Ansar and Daniilidis [48] for use in robotics applications at the Jet Propulsion Laboratory, California Institute of Technology [49,50]. As shown in Fig. 7, typically, approximately 2500–3500 pairwise matches are found for each image pair, resulting in 1000–1500 triplets over each set of 3 images. These matches are used for both camera pose estimation and for generating a high-speed 3D reconstruction of the parachute canopy as it is deployed.

As in Sec. IV.B, camera pose is corrected using a least-squares solver that minimizes reprojection error. Instead of using uniform grids of dots on planar surfaces, the pose correction step uses feature triplets found anywhere in the overlapping field of view of the three cameras. During flight, most of these triplets will be correspondences discovered on the canopy surface. The objective function triangulates these triplets and reprojects their intersection points back onto the image planes. The pose is optimal when it minimizes the residual error between the original and reprojected triplets.

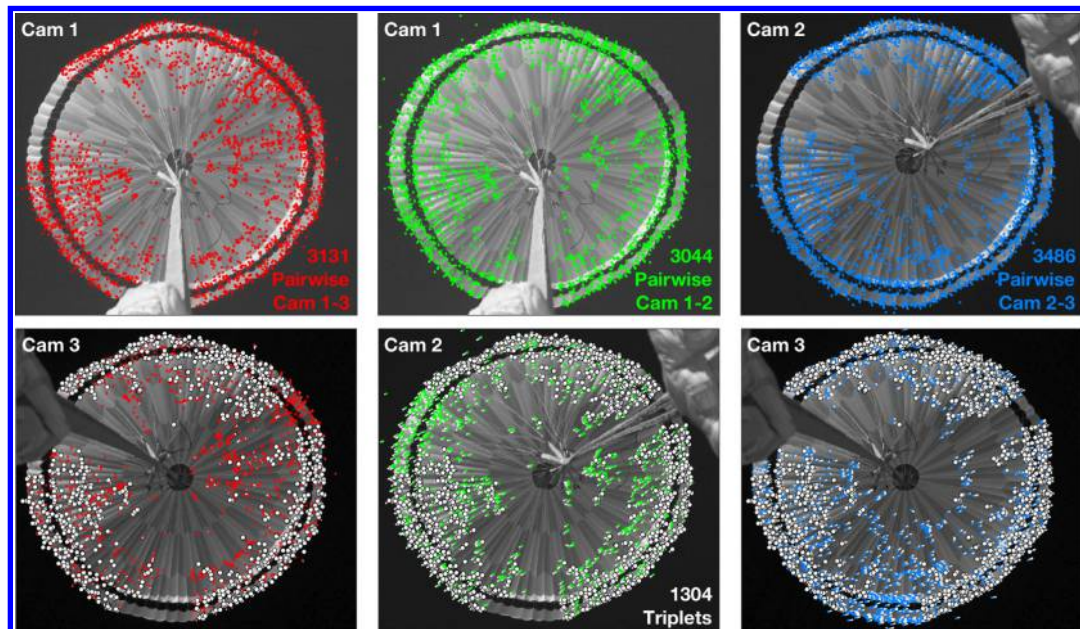
Unlike target-board dots, feature matches can occur anywhere in the scene at unspecified locations. Despite this lack of constraints, the minimization procedure still works in practice because of 1) the large number of points available for fitting, 2) the high quality of matches, and 3) the small number of fit parameters. Specifically, only six parameters are varied that are believed to be the most probable to change during flight, namely, the pitch, yaw, and roll of cameras 2 and 3 with respect to camera 1. High-quality matches are obtained by enforcing a consistency check over the three images: three pairwise matches form a triplet only if successive matches from camera 1 to camera 2, camera 2 to camera 3, and camera 3 to camera 1 lead back

to the exact same coordinate. A missing feature or an erroneous association between one or more of these image pairs is likely to cause this check to fail. In practice, less than half of pairwise matches are part of a triplet.

As the parachute canopy is visible for at least 4.8 s in each 1000 Hz image collection, roughly 5 million triplets are extracted during flight. Of these, 100,000 triplets are randomly chosen to generate a single set of corrected camera models for the flight. The magnitude of the final corrections in pitch, yaw, and roll is on the order of a few tenths of a degree, with the smallest changes occurring in roll. Figure 8 illustrates the importance of correcting the original camera models for small changes in pose. Before correction, poor alignment between separate camera pairs results in the large depth disparity observed in the left-most plot. By comparison, the corrected point clouds on the right show noticeably less surface variation. In the left-most plot of Fig. 8, pairwise matches (aggregated over all 3 baselines) are triangulated using camera models calibrated before launch. The distorted appearance seen is a consequence of small changes in camera pose that occurred after the original camera calibration. By minimizing the reprojection error of feature triplets, the pose is corrected, resulting in the center image in Fig. 8, which shows a pairwise surface reconstruction. The right image shows a more accurate but less complete reconstruction based on feature triplets.

To track the evolution of the parachute canopy over time, two separate point clouds are generated at 1 ms time intervals throughout the flight. These are the *pairwise*- and *triplet*-based point clouds plotted on the right side of Fig. 8.

Triplet-based clouds: As depth uncertainty typically decreases when the same point is viewed from multiple angles [51], the point cloud based on triplets is intended to provide the most accurate representation of the parachute canopy. This accuracy comes at the expense of coverage, however, because features that are occluded in any one of the three camera views will not be part of a triplet.

**Fig. 7** Example showing SIFT feature correspondences matched across a single set of synchronized images taken 1.500 s after mortar fire.

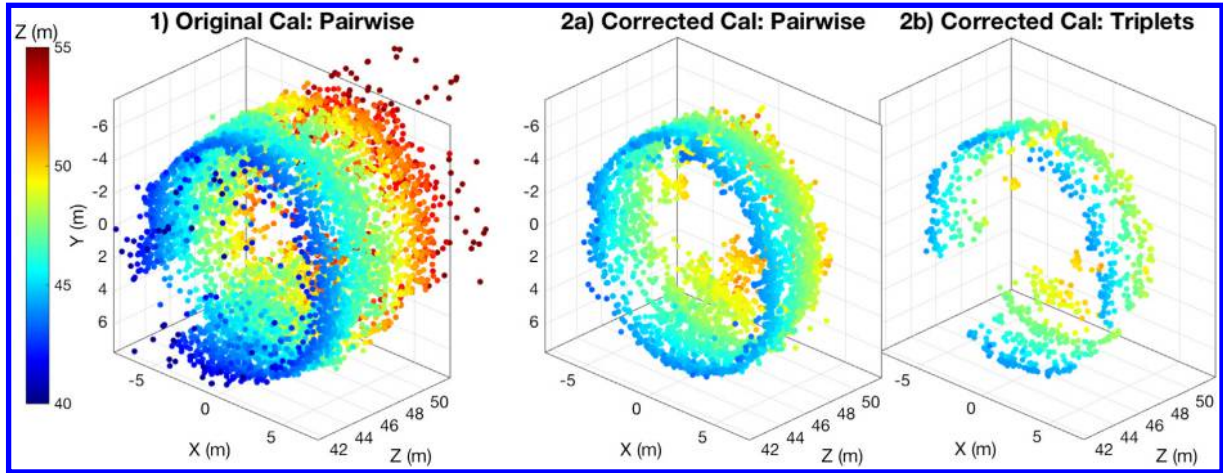


Fig. 8 Point clouds generated by triangulating the feature correspondences shown in Fig. 7. Each point cloud is color-coded by depth. 1) Pairwise matches using prelaunch calibration, 2a) pose-corrected pairwise surface reconstruction, and 2b) pose-corrected triplet matches.

Pairwise-based clouds: For applications where coverage is more important than accuracy, such as the projected area calculations performed in Sec. V.B, a second point cloud based on pairwise matches provides more complete coverage. This is especially true near the middle of the canopy, where lines converge and cause the greatest occlusion and confusion.

Based on localized gradient histograms, the SIFT features used in our analysis have a number of properties that help to address challenges discussed in Sec. IV.A. They provide limited invariance to lighting variations, affine transformations, and small changes in the appearance of surface features as seen from different camera angles. They are, however, not immune to occlusion from suspension lines and other parachute support structures, which can lead to spurious false matches. For this reason, the final point clouds generated in Fig. 6 are cleaned by backprojecting each point and computing its error with respect to the original matches. Typically 2–5% of points with errors greater than 0.5 pixels are removed from each cloud.

V. Data Products

This section outlines the different data products that are calculated based on the 3D point clouds generated during the ASPIRE tests.

Detailed data processing is only performed for the SR02 and SR03 flights, as only two of the three science cameras recorded data during the SR01 flight, whereas all three cameras recorded data for SR02 and SR03.

A. Vent Hole Tracking

The vent hole located in the middle of the parachute tracked using a disk-shaped template that matches the size and shape of the vent hole. As the size and appearance of the vent hole does not change significantly throughout each flight test, it can be followed without specialized fiducial markers despite intermittent occlusion from suspension lines and other hardware. Coordinates are estimated by 1) convolving each image with the template, and 2) interpolating the location of the maximum in the resulting convolved images [52]. These coordinates are then compared across synchronized image sets to triangulate the 3D trajectory of the vent hole during flight. The results of this process are shown in Fig. 9. Approximately 8–12% of coordinates are removed based on low match scores or high reprojection error, caused primarily by intermittent obstruction of the hole by suspension lines in one or more camera views.

The largest error in the trajectory is observed in the depth direction, where small errors in the estimated image coordinates lead to depth




Fig. 9 The trajectory of the central vent hole as calculated for SR02. Here the trajectory is plotted in the Camera Coordinate System, with its origin at the midpoint between three HS cameras oriented along the z axis.

errors that scale quadratically with distance [37]. For this reason, two vent-hole data products are generated: 1) the original triangulated coordinates, and 2) a corresponding set of distance-smoothed coordinates. The smoothing process involves converting the trajectory to spherical coordinates and replacing $r(t)$ with a best-fit third-order polynomial. This step removes transient errors in $r(t)$ that are likely to be artifacts of intermittent occlusion.

B. Projected Area Calculations

1. Overview

To calculate the projected area of the parachute, the previously computed point cloud based on pairwise matches with coordinates in the camera frame is used. To project the points into a two-dimensional (2D) cross section, the Z coordinate is dropped from the 3D points, such that the area is determined in a plane perpendicular to the line of sight of the cameras. The boundary of the 2D points is computed, and then the area is calculated from these ordered boundary points. To determine the boundary of the parachute, a convex hull of the points is generated. It is observed that the largest cross section of the parachute does not exhibit much concavity throughout flight, and the convex hull allows a relevant area to be calculated, even if the entire outer edge of the parachute is not resolved during the 3D reconstruction. Figure 10 illustrates the convex hull for a set of points, as well as a tighter boundary generated from a naive boundary fitting algorithm. If there are occluded regions in the point cloud, a more aggressive choice of the shrink factor in the algorithm can cause the algorithm to underestimate the projected area of the parachute, as the algorithm has no prior information or constraints on the shape of the parachute boundary. In practice, the area generated using the convex hull boundary differs from the area calculated with a tighter boundary only by approximately 1–2% and is therefore not a significant source of error in these calculations.

2. Possible Error Sources

As the method of computing area based on a polygon composed of 2D points is fairly straightforward, the error of this method can be estimated based on the uncertainty of the chosen boundary points. To perform this error estimation, the initial uncertainty for the XY coordinate of the point, currently $\sigma_{\text{point}} = 0.5$ m, is propagated through the Shoelace formula, which is the equation used to compute the area of the polygon formed by the points. This initial point uncertainty is estimated empirically in a separate procedure by selecting some points on the boundary of the canopy, sampling a large number of noisy points from a normal distribution with the image coordinate as the mean, and the standard deviation representing

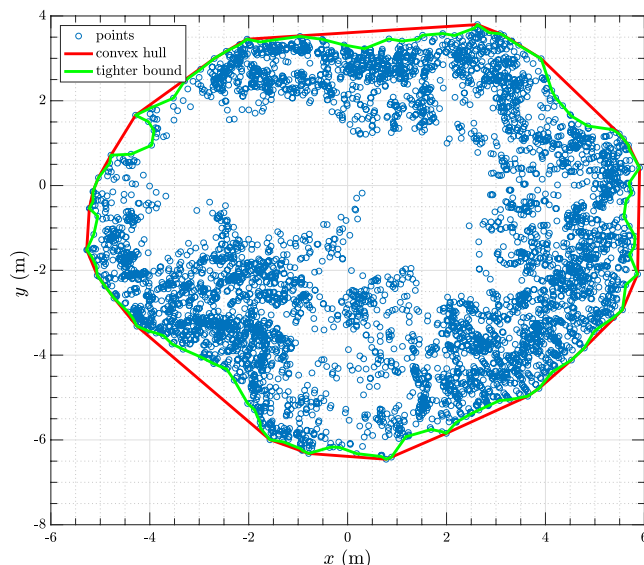


Fig. 10 Sample visualization of a projected area calculation where two different algorithms are used to calculate the parachute boundary.

the calibration and matching error, and then triangulating these noisy points and observing the calculated standard deviation. The final uncertainty for the area calculation is $\sigma_{\text{area}} \approx 1.5$ m², which can be as high as ~15% relative error for times early in the inflation process, but is ~1% after full inflation.

Although the baseline uncertainty from the error in triangulation of 3D points is small, it is believed that the main contribution to the error when estimating the projected area is whether the boundary points that are selected accurately represent the largest cross section of the parachute. The primary source of error that affects whether these points can be captured is occlusions of the parachute skirt. Though one major source of occlusion is the confluence fitting and riser, this is mostly mitigated by having a third camera, which fills in the occlusions of the two other cameras. Another potential source of occlusion of the parachute is self-occlusion caused by motion of the parachute away from the vehicle center line. For simplicity, in this work, tilt angle will be used to refer to the angular offset between the vehicle center line and the parachute axis of symmetry. Overlaying the projected area with the tilt angle computed from the vent hole location for SR02 (Fig. 11), there appears to be a correlation between a larger rotation angle of the parachute and a drop in calculated projected area. This drop in area is believed to be due to occlusion of aspects of the canopy, as the analytical decrease in the area due to the rotation does not account for the large calculated drops, which are observed to be up to 8%. A possible reason as to why this effect is observed is that rotation of the parachute can cause the widest part of the skirt to be occluded, which prevents the reconstruction algorithm from identifying this part of the parachute. This phenomenon is significant for SR02, where there is a relatively large parachute motion, but is not as significant for SR03, which has minimal solid-body motion. This is believed to contribute to the relatively smooth (fewer oscillations) calculated projected area throughout the flight for SR03 compared with SR02. Though it would be possible to project the reconstructed points onto a plane that is perpendicular to a vector running from the cameras to the vent hole, this is not performed because it does not remove the error associated with camera occlusions, and would only account for a possible solid-body rotation of the parachute.

C. Parachute Canopy Velocity

Two methods are used to compute and compare local canopy velocity measurements: an image-based method using feature points, and a method using the 3D coordinates of point clouds directly. For both methods, the points used are prethresholded by distance to attempt to capture points on the parachute skirt only. Figure 12 illustrates the results of computing velocities for points of interest from the two different methods for the same image frame. These velocities




Fig. 11 Tilt angle (calculated from smoothed vent hole coordinates) compared with area for SR02.

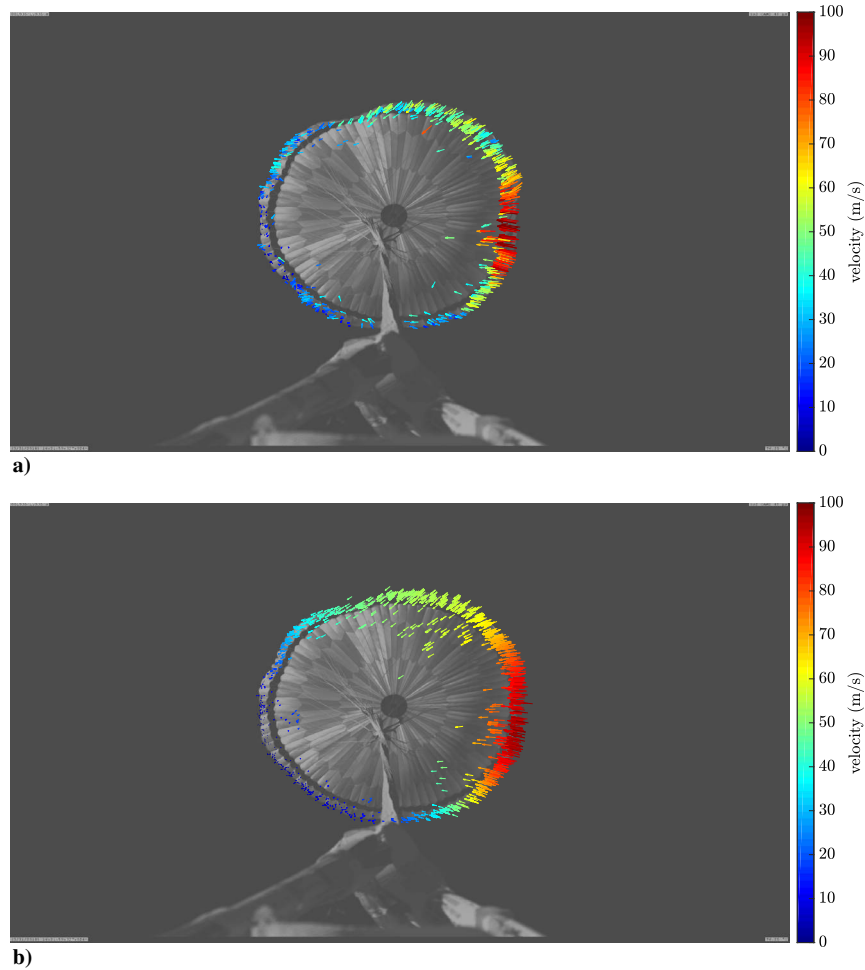


Fig. 12 Two different methods of computing in-plane canopy velocity: a) image feature tracking then rescaled to calculate velocity, and b) point set alignment.

are computed in-plane (X – Y) as visualized in the relevant figures, and represent a 2D velocity.

In the image-based approach, feature points that are used to generate the point clouds are selected to be tracked across successive frames. The points are tracked using the Kanade–Lucas–Tomasi (KLT) feature tracking algorithm [53]. After tracking, a pixel motion is obtained for each point, which is then rescaled into a dimensional velocity (m/s) using the following:

$$v_{\text{dimensional}} = \frac{v_{\text{pixel}} \times Z}{f} \times \text{frame_rate} \quad (1)$$

where v_{pixel} is measured in pixels per frame, Z is the depth of the feature point in meters, f is the focal length of the camera in pixels, and frame_rate is the frame rate, which is equal to 1000 frames per second for this work. The main benefit of this method is that estimating the motion in image space is accurate. However, this method is challenging to apply when large motions are observed from frame to frame, i.e. when large parachute motion and deformations will inhibit feature tracking.

In the direct approach, which uses the 3D reconstructed points, the goal is to compute a motion for each point that transforms the current point cloud into the next point cloud. The pointwise motion will correspond to the velocity of the canopy points. The method of coherent point drift [54] is used, which performs nonrigid alignment of unorganized point clouds. This method represents the point sets using a Gaussian mixture model. The benefit of representing the data with a Gaussian mixture model is that the correspondence problem between point clouds does not have to be solved, as noisy

data are handled naturally, and therefore the computed motion is smooth. Points that are near each other should have similar velocities. In contrast, with the image-based method, the measurements generated from this point set alignment may be less accurate because the inputs are noisy triangulated points and contain outliers. Image feature tracking generally works well only with small motions and the assumption of constant brightness of the feature points, which can limit this method to small frame differences for the current application. However, the point set alignment method can potentially compute motion across 10 s of frames because the inputs are simply point clouds, as long as the shapes are not dramatically different. Future work could involve smoothing the results over a larger range to overcome the single-frame inaccuracies. Points are aligned in 2D to decrease processing time and focus on velocity components in the X – Y plane.

As an optional postprocessing step, results are also generated with the vent hole velocity subtracted. This is useful for the postinflation phase, where large solid-body parachute motion can dominate the observed motion. Figure 13 illustrates an instance in the trajectory where the dominant motion is a rigid body mode by the entire canopy, as opposed to local relative motion. After subtracting the vent velocity, the actual velocity of the feature points is nearly zero.

VI. Discussion

Although the majority of the current effort has focused on generating the previously described data products, some preliminary work has been completed on analyzing these results to gain better insight into the supersonic parachute dynamics observed during the ASPIRE flight tests.

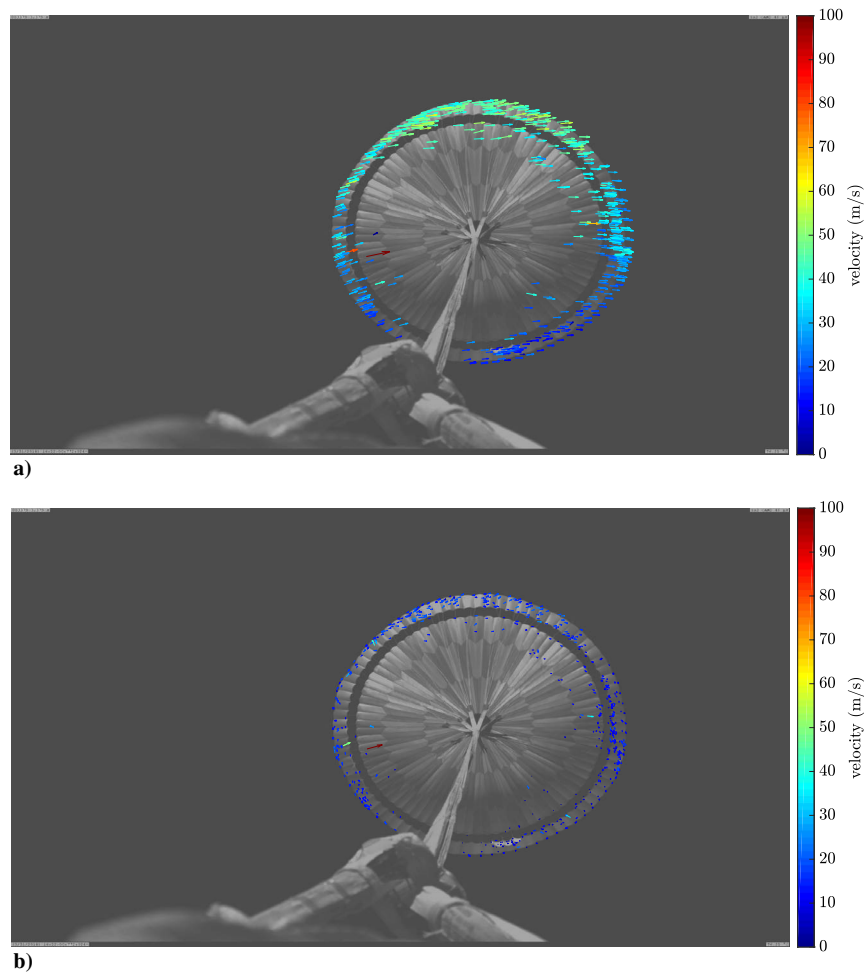


Fig. 13 Canopy in-plane velocity without (a) and with (b) subtracting the vent hole smoothed velocity.

A. Load and Projected Area Data

The calculated projected areas of the parachutes for SR02 and SR03 are shown in Fig. 14, and the comparison of the normalized projected area and measured load is shown in Fig. 15. A general trend seen in the projected area data is that the projected area tends to

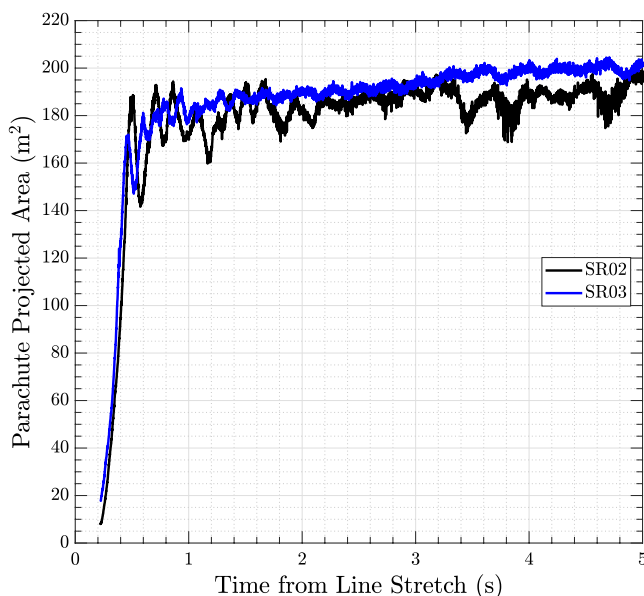


Fig. 14 Projected area versus time for SR02 and SR03.

lag behind the peaks in measured load. Furthermore, as the measured load decreases postinflation for both SR02 and SR03, a small increase in projected area is observed.

B. Inflation Dynamics

The same numerical method for point alignment that is used to compute the parachute canopy velocity can also be leveraged to fit a contour to desired points in a 3D point cloud. Fitting contours to the points that comprise the skirt of the parachute can provide some insight into the dynamics that occur during canopy inflation. Figure 16 shows a calculated contour that has been fitted to points found to be in the parachute skirt, and then the points are reprojected back onto the original image in order to visualize how accurately the skirt is captured. Figure 17 shows fitted canopy skirt contours color-coded by time for SR02 and SR03. The times shown span the inflation process up to the point where peak load is measured. From Sec. V.C, it can be seen that parts of the canopy skirt are moving at ~ 100 m/s during inflation.

C. Inflated Canopy Shape

Figure 18 shows two sample point clouds from the SR02 flight test. As previously discussed, the discrepancy between the point cloud generated with pairwise matches versus triplets is clearly seen, with the triplet matches resolving less of the canopy surface. However, it is believed that the error associated with the XYZ coordinates of the parachute surface generated with the triplet matches is lower than the error associated with the points with pairwise matches.

Figures 19 and 20 show variations of the SR02 and SR03 canopy profiles postinflation. Colors correspond to the probability that the reconstructed canopy shape was located in a specific location. In addition, the experimental data are compared with an Finite Element

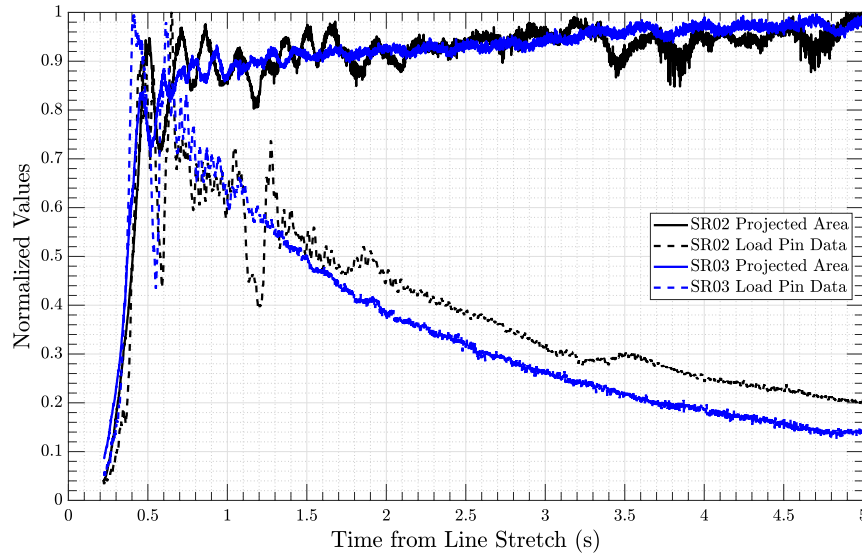


Fig. 15 Normalized load and projected area for SR02 and SR03 flights. The load and projected area are normalized by each of their respective maximum values.

Analysis (FEA)-derived symmetric DGB profile. Contours in Figs. 19 and 20 represent a normalized Probability Density Function (PDF) of Z coordinates of the canopy location for a given r , where r is defined to be $r = ((X - x_c)^2 + (Y - y_c)^2)^{0.5}$, where x_c and y_c are the X and Y locations of the center of the vent hole (for a given frame). This allows a side profile view of the canopy to be visualized, similar to profiles presented in [38], and all reconstructed features can be collapsed into a single two-dimensional plot. These distributions of Z coordinates are






Fig. 16 Estimated skirt contour overlaid on the original parachute image frame for a) SR02 and b) SR03.



a)



b)

Fig. 17 Fitted contours to the skirt points color-coded by time from line stretch for a) SR02 and b) SR03.

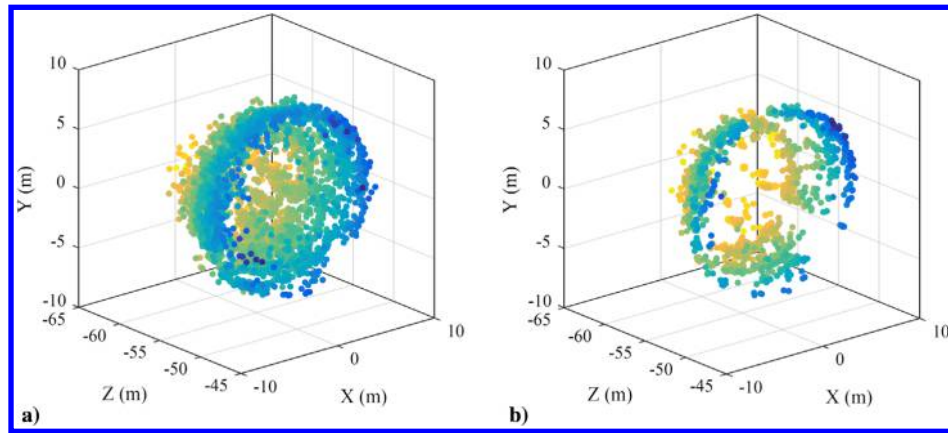


Fig. 18 Sample point clouds from the SR02 flight. a) Pairwise matches, and b) triplet matches. The point clouds show the location of features on the canopy 1.89 s after mortar fire. Features have been colored by their Z coordinate.

created by using all triplet matches in the SR02 and SR03 images starting ~ 1 s after parachute deployment. Distributions are normalized on a per r -value basis in order to account for the fact that the triplet point clouds generally find more points near the skirt compared with the center of the parachute, as previously discussed. Without normalization, the structure of the canopy closer to the vent hole is more difficult to visualize.

Figures 19 and 20 also include an FEA-derived DGB profile as a point of comparison (aligned qualitatively). Although the experimental data are quite noisy, there do seem to be some curvature differences between the experimental data and the FEA-derived DGB

profile. As the parachute inflation conditions were different between the SR02 and SR03 tests, it is expected that the inflated profiles of the canopies will also differ. The FEA-derived profile is included as a qualitative point of comparison only, and the same profile is shown in Figs. 19 and 20, even though the flow conditions differ between the experiments. Furthermore, the parachute canopy shape is expected to change as a function of time, as the velocity of the vehicle and the atmospheric conditions change during descent of the payload. However, this qualitative comparison does highlight the fact that further investigation is required in order to validate the assumed shape of canopy profiles that are used in numerical simulations.




Fig. 19 Reconstructed canopy shape probability density function for SR02 (contours) compared with an FEA-derived DGB profile (markers).

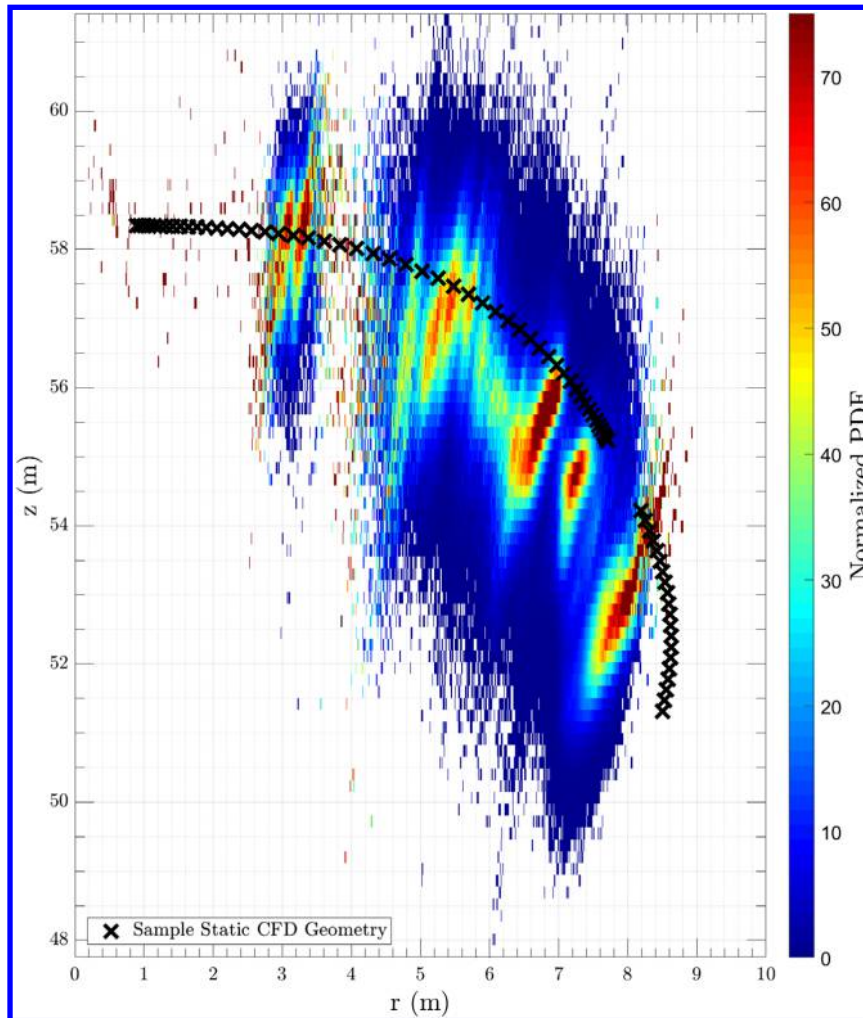


Fig. 20 Reconstructed canopy shape probability density function for SR03 (contours) compared with an FEA-derived DGB profile (markers).

VII. Conclusions

This work has provided an overview of the ASPIRE test campaign and image processing techniques used to create time-resolved point clouds representing the parachute surface as a function of time. These point clouds were then used to further look into different physical phenomena associated with supersonic parachute dynamics. Due to a variety of error sources presented in this work, the 3D reconstructed point clouds can be noisy, and preliminary work has been completed to compare the experimentally observed canopy shape with FEA-derived axis-symmetric models.

Acknowledgments

The research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with NASA. The authors would like to acknowledge the contributions of the entire Advanced Supersonic Parachute Inflation Research Experiment (ASPIRE) team and would also like to thank Yang Cheng and Adnan Ansar for providing software and advice.

References

- [1] Cooley, C. G., and Lewis, J. G., "Viking 75 Project: Viking Lander System Primary Mission Performance Report," CR-145148, April 1977.
- [2] Eckstrom, C. V., and Preisser, J. S., "Flight Test of a 30-Foot-Nominal-Diameter Disk-Gap-Band Parachute Deployed at a Mach Number of 1.56 and a Dynamic Pressure of 11.4 Pounds Per Square Foot," NASA TM X-1451, Aug. 1967.
- [3] Preisser, J. S., and Eckstrom, C. V., "Flight Test of a 40-Foot-Nominal-Diameter Disk-Gap-Band Parachute Deployed at a Mach Number of 1.91 and a Dynamic Pressure of 11.6 Pounds Per Square Foot," NASA TM X-1575, Aug. 1968.
- [4] Preisser, J. S., and Grow, R. B., "High Altitude Flight Test of a Reefed 12.2-Meter Diameter Disk-Gap-Band Parachute With Deployment at a Mach Number of 2.58," NASA TN D-6469, Aug. 1971.
- [5] Fallon, E. J., "System Design Overview of the Mars Pathfinder Parachute Decelerator Subsystem," AIAA Paper 1997-1511, 1997.
- [6] Spencer, D. A., Blanchard, R., Braun, R. D., Kallemeyn, P. H., and Thurman, S. W., "Mars Pathfinder Entry, Descent, and Landing Reconstruction," *Journal of Spacecraft and Rockets*, Vol. 36, No. 3, 1999, pp. 357–366.
<https://doi.org/10.2514/2.3478>
- [7] Cruz, J. R., and Lingard, J. S., "Aerodynamic Decelerators for Planetary Exploration: Past, Present, and Future," AIAA Paper 2006-6792, 2006.
- [8] Witkowski, A., and Bruno, R., "Mars Exploration Rover Parachute Decelerator System Program Overview," AIAA Paper 2003-2100, 2003.
- [9] Adams, D. S., Witkowski, A., and Kandis, M., "Phoenix Mars Scout Parachute Flight Behavior and Observations," *2011 IEEE Aerospace Conference*, IEEE, New York, 2011, pp. 1–8.
- [10] Cruz, J. R., Way, D. W., Shidner, J. D., Davis, J. L., Adams, D. S., and Kipp, D. S., "Reconstruction of the Mars Science Laboratory Performance and Comparison to the Descent Simulation," AIAA Paper 2013-1250, 2013.
- [11] Kipp, D. M., and Buecher, D., "A Brief History of InSight Parachute Development and Acceptance for Flight," *15th International Planetary Probes Workshop*, Boulder, CO, June 2018.
- [12] Clark, I., and Tanner, C., "A Historical Summary of the Design, Development, and Analysis of the Disk-Gap-Band Parachute," *IEEE Aerospace Conference*, Paper 2466, Big Sky, MT, 2017.
- [13] O'Farrell, C., Karlgaard, C., Tynis, J. A., and Clark, I. G., "Overview and Reconstruction of the ASPIRE Project's SR01 Supersonic

- Parachute Test," *2018 IEEE Aerospace Conference*, Inst. of Electrical and Electronics Engineers, New York, 2018, pp. 1–18.
- [14] O'Farrell, C., Sonneveldt, B. S., Karlgaard, C., Tynis, J. A., and Clark, I. G., "Overview of the ASPIRE Project's Supersonic Flight Tests of a Strengthened DGB Parachute," *2019 IEEE Aerospace Conference*, Inst. of Electrical and Electronics Engineers, New York, 2019, pp. 1–18.
 - [15] Bendura, R. J., Huckins, E. K., and Coltrane, L. C., "Performance of a 19.7-Meter-Diameter Disk-Gap-Band Parachute in a Simulated Martian Environment," NASA TM X-1499, Aug. 1968.
 - [16] Eckstrom, C. V., and Preisser, J. S., "Flight Test of a 40-Foot-Nominal-Diameter Disk-Gap-Band Parachute Deployed at a Mach Number of 2.72 and a Dynamic Pressure of 9.72 Pounds per Square Foot," NASA TM X-1623, Aug. 1968.
 - [17] Eckstrom, C. V., "Flight Test of a 40-Foot-Nominal-Diameter Disk-Gap-Band Parachute Deployed at a Mach Number of 3.31 and a Dynamic Pressure of 10.6 Pounds per Square Foot," NASA TM X-1924, Feb. 1970.
 - [18] Eckstrom, C. V., and Branscome, D. R., "High-Altitude Flight Test of a Disk-Gap-Band Parachute Deployed Behind a Bluff Body at a Mach Number of 2.69," NASA TM X-2671, Dec. 1972.
 - [19] Dickinson, D., Hicks, F., Schlemmer, J., Michel, F., and Moog, R. D., "Balloon Launched Decelerator Test Program Post-Flight Test Report: BLDT Vehicle AV-1," NASA CR-112176, 1972.
 - [20] Dickinson, D., Hicks, F., Schlemmer, J., Michel, F., and Moog, R. D., "Balloon Launched Decelerator Test Program Post-Flight Test Report: BLDT Vehicle AV-2," NASA CR-112177, 1972.
 - [21] Dickinson, D., Hicks, F., Schlemmer, J., Michel, F., and Moog, R. D., "Balloon Launched Decelerator Test Program Post-Flight Test Report: BLDT Vehicle AV-4," NASA CR-112179, 1972.
 - [22] Clark, I. G., Adler, M., and Manning, R., "Summary of the First High-Altitude, Supersonic Flight Dynamics Test for the Low Density Supersonic Decelerator Project," AIAA Paper 2015-2100, 2015.
 - [23] Ray, E. S., Bretz, D. R., and Morris, A., "Photogrammetric Analysis of CPAS Main Parachutes," AIAA Paper 2011-2538, 2011.
 - [24] Ray, E. S., "Photographic Volume Estimation of CPAS Main Parachutes," AIAA Paper 2017-3229, 2017.
 - [25] Brent, R. J., and DeWeese, J. H., "Filling Time Prediction Approaches for Solid-Cloth Type Parachutes," AIAA Paper 1966-1503, 1966.
 - [26] Houmard, J., "Stress Analysis of the Viking Parachute," AIAA Paper 1973-444, 1973.
 - [27] Schoenenberger, M., Queen, E., and Cruz, J. R., "Parachute Aerodynamics from Video Data," AIAA Paper 2005-1633, 2005.
 - [28] Underwood, J. C., Saunders, A., Lingard, J. S., Marraffa, L., Ferracina, L., Larosekand, G. L., and McTavish, S. A., "Subsonic Wind Tunnel Testing of Various Parachute Types," AIAA Paper 2015-2112, 2015.
 - [29] Cruz, J. R., O'Farrell, C., Zumwalt, C. H., and Keller, D. F., "Wind Tunnel Testing of Ringsail and Disk-Gap-Band Parachutes," NASA Engineering and Safety Center Tech. Rept. 14-00932, 2015.
 - [30] Mayhue, R. J., "Drag Characteristics of a Disk-Gap-Band Parachute with a Nominal Diameter of 1.65 Meters at Mach Numbers From 2.0 to 3.0," NASA Langley Research Center TN D-6894, Dec. 1972.
 - [31] Schairer, E. T., Kushner, L. K., Heineck, J. T., and Solis, E., "Measurements of Parachute Dynamics in the World's Largest Wind Tunnel by Stereo Photogrammetry," AIAA Paper 2018-3802, 2018.
 - [32] Sengupta, A., "Fluid Structure Interaction of Parachutes in Supersonic Planetary Entry," AIAA Paper 2011-2541, 2011.
 - [33] Tanner, C. L., Clark, I. G., and Chen, A., "Overview of the Mars 2020 Parachute Risk Reduction Plan," *IEEE Aerospace Conference*, IEEE, New York, 2018.
 - [34] Karlgaard, C., Tynis, J. A., and O'Farrell, C., "Reconstruction of the Advanced Supersonic Parachute Inflation Research Experiment Sounding Rocket Flight Test," AIAA Paper 2018-3624, 2018.
 - [35] Karlgaard, C., Tynis, J. A., O'Farrell, C., and Sonneveldt, B., "Reconstruction of the Advanced Supersonic Parachute Inflation Research Experiment Sounding Rocket Flight Tests with Strengthened Disk-Gap-Band Parachute," *AIAA Scitech 2019 Forum*, AIAA Paper 2019-0014, Jan. 2019.
 - [36] Wöhler, C., and d'Angelo, P., "Stereo Image Analysis of Non-Lambertian Surfaces," *International Journal of Computer Vision*, Vol. 81, No. 2, 2008, pp. 172–190.
<https://doi.org/10.1007/s11263-008-0157-1>
 - [37] Gallup, D., Frahm, J., Mordohai, P., and Pollefeys, M., "Variable Baseline/Resolution Stereo," *2008 IEEE Conference on Computer Vision and Pattern Recognition*, IEEE, New York, 2008, pp. 1–8.
<https://doi.org/10.1109/CVPR.2008.4587671>
 - [38] Schairer, E. T., Kushner, L. K., Heineck, J. T., and Solis, E., "Measurements of Parachute Dynamics in the World's Largest Wind Tunnel by Stereo Photogrammetry," *2018 Aerodynamic Measurement Technology and Ground Testing Conference*, AIAA Paper 2018-3802, June 2018.
<https://doi.org/10.2514/6.2018-3802>
 - [39] Weng, J., Cohen, P., and Herniou, M., "Camera Calibration with Distortion Models and Accuracy Evaluation," *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, Vol. 14, No. 10, 1992, pp. 965–980.
<https://doi.org/10.1109/34.159901>
 - [40] Zhang, Z., "A Flexible New Technique for Camera Calibration," *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, Vol. 22, No. 11, 2000, pp. 1330–1334.
<https://doi.org/10.1109/34.888718>
 - [41] Roth, G., "Computing Camera Positions from a Multi-Camera Head," *Third International Conference on 3-D Digital Imaging and Modeling*, IEEE Computer Soc., New York, 2001, pp. 135–142.
 - [42] Heikkilä, J., and Silven, O., "A Four-Step Camera Calibration Procedure with Implicit Image Correction," *Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, IEEE, New York, 1997, pp. 1106–1112.
 - [43] Di, K., and Li, R., "CAHVOR Camera Model and Its Photogrammetric Conversion for Planetary Applications—Di—2004—Journal of Geophysical Research: Planets (1991–2012)—Wiley Online Library," *Journal of Geophysical Research*, Vol. 109, No. E4, 2004, Paper E04004.
<https://doi.org/10.1029/2003JE002199>
 - [44] Kim, W. S., Ansar, A. I., and Steele, R. D., "Rover Mast Calibration, Exact Camera Pointing, and Camera Handoff for Visual Target Tracking," *ICAR '05. Proceedings of 12th International Conference on Advanced Robotics*, Seattle, WA, 2005, pp. 384–391.
<https://doi.org/10.1109/ICAR.2005.1507439>
 - [45] Lowe, D. G., "Distinctive Image Features from Scale-Invariant Key-points," *International Journal of Computer Vision*, Vol. 60, No. 2, 2004, pp. 91–110.
<https://doi.org/10.1023/B:VISI.0000029664.99615.94>
 - [46] Vedaldi, A., and Fulkerson, B., *Vlfeat: An Open and Portable Library of Computer Vision Algorithms*, Assoc. for Computing Machinery, New York, 2008, <http://www.vlfeat.org/>.
 - [47] Yousif, K., Bab-Hadiashar, A., and Hoseinnezhad, R., "An Overview to Visual Odometry and Visual SLAM: Applications to Mobile Robotics," *Intelligent Industrial Systems*, Vol. 1, No. 4, 2015, pp. 289–311.
<https://doi.org/10.1007/s40903-015-0032-7>
 - [48] Ansar, A., and Daniilidis, K., "Linear Pose Estimation from Points or Lines," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 25, No. 5, 2003, pp. 578–589.
<https://doi.org/10.1109/TPAMI.2003.1195992>
 - [49] Clouse, D., Padgett, C., Ansar, A., and Cheng, Y., "Automated Camera Array Fine Calibration," NASA Tech Brief, NPO-45505, Jet Propulsion Lab., California Inst. of Technology, Pasadena, CA, 2008.
 - [50] Johnson, A., Ansar, A., Matthies, L., Trawny, N., Mourikis, A., and Roumeliotis, S., "A General Approach to Terrain Relative Navigation for Planetary Landing," *AIAA Infotech@Aerospace 2007 Conference and Exhibit*, AIAA, Reston, VA, 2012, p. 391.
 - [51] Ruml, M., Irsch, A., and Bischof, H., "Multi-View Stereo: Redundancy Benefits for 3D Reconstruction," *35th Workshop of the Austrian Association for Pattern Recognition*, Vol. 4, 2011.
 - [52] Brunelli, R., *Template Matching Techniques in Computer Vision, Theory and Practice*, Wiley, Chichester, England, U.K., 2009, Chap. 3.
 - [53] Tomasi, C., and Kanade, T., "Detection and Tracking of Point Features," Carnegie Mellon Univ. Tech. Rept. CMU-CS-91-132, Pittsburgh, 1991.
 - [54] Myronenko, A., and Song, X. B., "Point-Set Registration: Coherent Point Drift," *CoRR*, Vol. abs/0905.2635, 2009, <http://arxiv.org/abs/0905.2635>.

J. A. Christian
Associate Editor