Unsteady Pressure Sensitive Paint Camera Calibration Improvements

Nicholas Califano[1]
Metis Technology Solutions, Moffett Field, CA, 94035

Nettie H. Roozeboom[2]
NASA Ames Research Center, Moffett Field, CA 94035

Data collection using unsteady pressure-sensitive paint (uPSP) has been greatly increased in recent large-scale demonstrations. With this significant increase comes challenges to process this magnitude of data. Techniques designed for several thousands of images collected in a lab do not necessarily scale to millions collected in a production-level system. Novel techniques are needed to meet the tighter requirements imposed on computation time, robustness, and accuracy. This paper outlines a number of such techniques and improvements in regards to the camera calibration process.

I. Introduction

Over the past two decades pressure-sensitive paint (PSP)and in particular uPSP has become a powerful technique to measure surface pressure of a wind tunnel model with high spatial density [1–5]. In the past, external calibration, the process by which the relative position and orientation between the camera and model is found, was previously performed for only one frame of a run sequence. This was sufficient because either because only one frame was used [6], or model motion between frames was assumed to be small [4]. However, with the development of unsteady pressure-sensitive paint and use of large wind tunnel models at high wind speed comes stricter requirements.

In the most recent large-scale wind tunnel tests with uPSP at the NASA Ames Unitary Plan Wind Tunnel (UPWT) [2,4], data was collected for the Space Launch System (SLS). These tests presented challenges as hundreds of run sequences were collected each with tens of thousands of frames. The large model necessitated a large field of view and cause significant sting deflection and model oscillation for high Mach number run sequences. The model motion between successive frames was originally handled with image registration in order to map the successive frames onto the first frame. This method however is slow, inaccurate, and does not lend itself to additional useful data products such as uncertainty quantification, model motion tracking over time, or node-to-pixel mapping for successive frames. Due to the magnitude of the problem several improvements were made to the camera calibration process in regards to the robustness, accuracy, and computational expense. These improvements will additionally lend themselves to even more dynamic models such as those with large deformation or even flutter.

Improvements were made to the internal calibration pipeline in regard to the data collection procedure and software capabilities. Improvements to the external calibration pipeline relate to the overall robustness, data collection procedure, processing pipeline, and overall computational expense.

II. Internal Calibration Improvements

A. Data Collection Procedure

Previously PSP utilized a large and cumbersome board for calibration. Additionally, the calibration procedure was poorly defined and lead to a relatively large uncertainty. Using Monte Carlo uncertainty analysis, the less-than-ideal internal calibration procedure led to an average uncertainty in the model external calibration of roughly 0.4° and 0.85”

[1] Software Engineer, Experimental Aero-Physics Branch, M.S. 260-1, AIAA Member
[2] Aerospace Engineer, Experimental Aero-Physics Branch, M.S. 260-1, AIAA Member

(on a distance of roughly 75”). In terms of pixel uncertainty, the average node-to-pixel mapping error is on the order of 0.3 pixels. A heatmap of the error on the SLS model is given in Fig. 1.

To improve the operational procedure, the board was updated to a smaller more manageable board. The new board is much flatter and stiffer than the previous board, is half the weight, and is easily held and manipulated by one person. Analysis has shown that by capturing images of the board in many positions in the tunnel, the calibration is more accurate even with the smaller board. The new procedure does not significantly increase the overall setup time in the wind tunnel. The new calibration procedure reduces the external calibration uncertainty to 0.08° and 0.16”. In terms of pixels, the new procedure has the same distribution, reduces the magnitude of the average node-to-pixel uncertainty to 0.06 pixels or less, as shown in Fig. 2. This process can also be easily scaled up or down to accommodate a lab setup or other production-level wind tunnel.

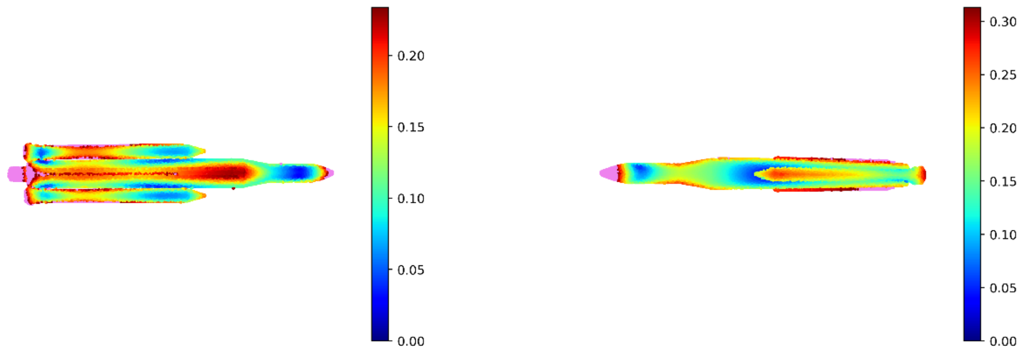

Fig. 2: Heatmap of Node-to-Pixel Mapping Uncertainty (in Pixels) for Old Procedure Pink Represents Oversaturated Error

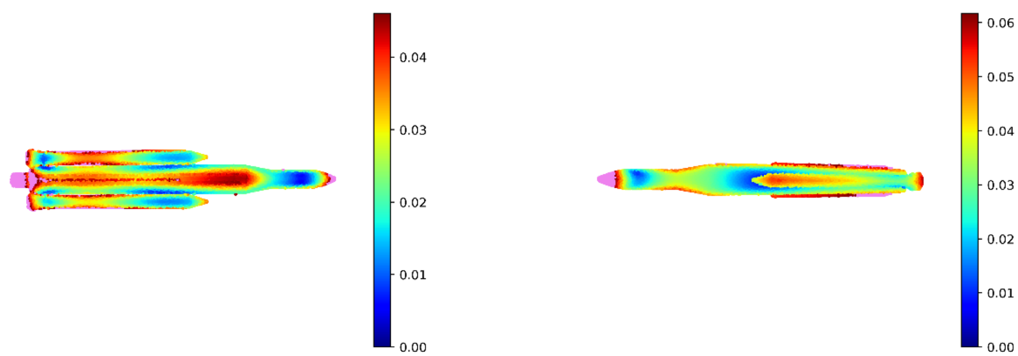

Fig. 1: Heatmap of Node-to-Pixel Mapping Uncertainty (in Pixels) for New Procedure Pink Represents Oversaturated Error

B. Software Capabilities

For internal calibration data processing, it is recommended to use Calib.io’s Calibrator software [7] due to its robustness, ease of use, visualizations, and uncertainty quantification. The original PSP internal calibration software allowed for occlusions but required significant manual input to process properly. Opensource software such as OpenCV [8] is available, but tunnel images are typically cluttered with background artifacts and tuning a blob detector or corner finder to those images also requires significant manual input. The Calib.io software is extremely robust to background objects and required little to no manual input.

The Calib,io software also provides several visualizations as to the quality of the internal calibration. These visualizations are useful for identifying the occasional frame with a mis-localized target, as well as more global biases such as regions missing in the calibration volume. The Calib.io software also provides uncertainty estimates for the calibration which are useful to estimate the overall external calibration uncertainty and provide a quantitative result for the internal calibration quality.

III. External Calibration Improvements

A. **Robustness**

1. Occlusion Checking

In order to associate the 3D model surface with the 2D images, the 3D model geometry is discretized into a surface mesh and every mesh node is projected into the image. An immediate problem occurs as many of the nodes are occluded by other nodes and should therefore not be associated with any pixel in the image. Previous iterations of PSP processing software utilized a z-buffer to reject occluded nodes [9,10], but that method has a disadvantage in that at most one node can be associated with a pixel. To avoid unassociated nodes, the node size has been tuned to be approximately the same size as one pixel. As geometry become more complex, the model becomes larger, and more cameras are utilized, sizing nodes in this way becomes difficult or impossible. Additionally, sizing the mesh nodes such that they are roughly equal to one pixel causes aliasing-like effects as show in Fig. 3 where not all pixels are utilized. As systems become larger, those unutilized pixels are an easy avenue to extract more data and information from a PSP system.

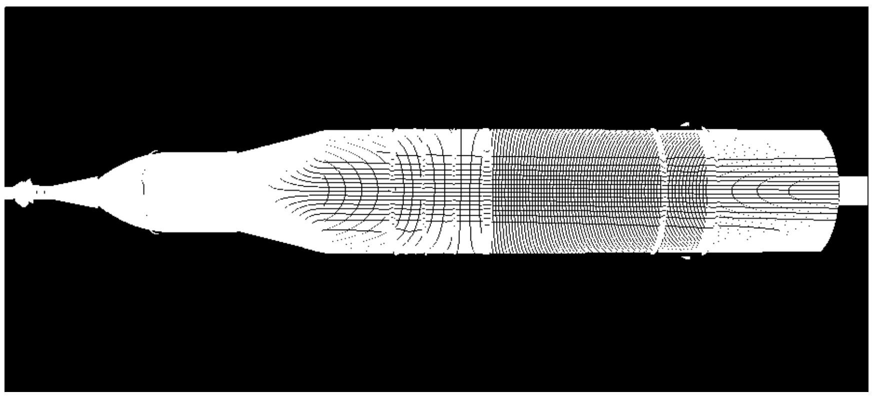

Fig. 3: Previous SLS tests visualization of pixels with at least one node

To check for occlusions, the new iteration of uPSP processing performs a back-face culling operation then utilizes a bounding volume hierarchy (BVH) for occlusions. The BVH occlusion checking operation is computationally expensive compared to the z-buffer but can associate an arbitrary number of nodes with a pixel. BVHs are the standard for ray tracing-based rendering and can have real-time performance [11]. The back-face culling is performed first since it is computationally inexpensive and can additionally reject nodes based on viewing angle. It is common in PSP applications to reject the visibility of nodes with a large viewing angle [4,12] because there are too few pixels imaging too many nodes, so combining these steps is computationally efficient.

2. Target Matching

The recent tests with the SLS presented several challenges to the traditional PSP external calibration processing pipeline. Typically, the images are well lit and have relatively uniform illumination both within an image and between images. Additionally, the models used are typically relatively small and have relatively simple geometry. However, the most recent test had a relatively large model and collected data from many run sequences. This presented challenges since the large model led to more instances of damaged paint, and the large number of run sequences led to more degradation of the paint from the increased UV exposure. Damaged paint, shadows, and complex model geometry led to many artifacts that were visually similar to the targets, and the paint degradation led to low contrast between the paint and the registration targets. All of these factors made it difficult to detect targets as well as ambiguous to match them. Several robustness measures were implemented to ensure the external calibration quality did not suffer.

A pair of representative images are shown in Fig. 4. Notice the contrast is high for the image taken at the start of the test but is low at the end of the test. Additionally, a different upper section was used during the second half of the test, and so the contrast between the paint and targets was different within an image for the upper and lower section of the model.

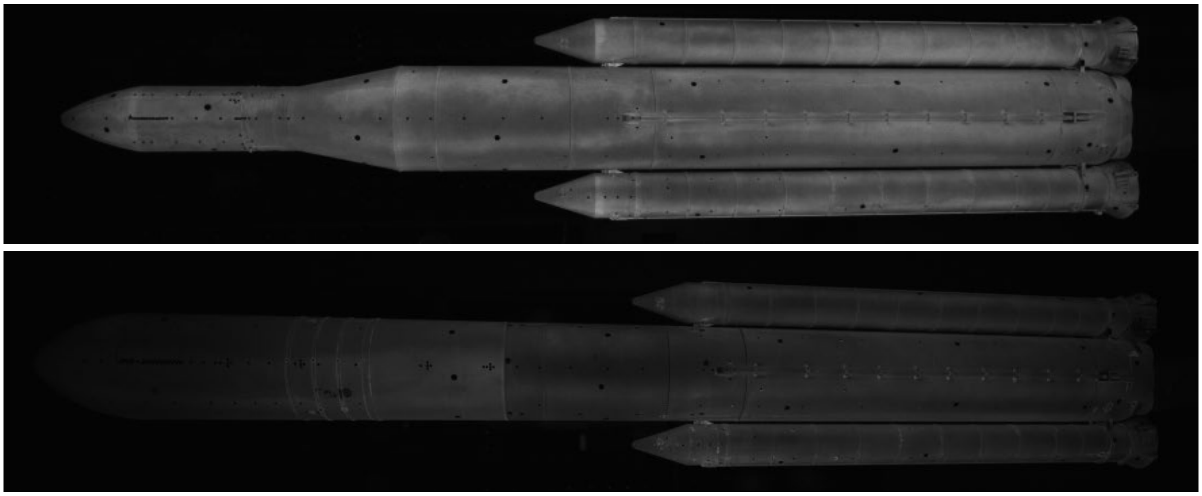

Fig. 4: Representative Images of Model at Start (Top) and End (Bottom) of Test

The robustness measures implemented were a number of checks on the matching scheme to ensure consistency properties. The first measure is a check that no two projected locations are too close to each other. If two 3D targets are projected within the minimum threshold distance, they are both rejected. The next measure is referred to as the bi-filter. The bi-filter first checks that all 3D targets have at most 1 image target within the maximum matching distance and visa versa. Any 3D targets or image targets that violate this are rejected. The final matching robustness measure is to reject targets that are on the visual boundary between two surfaces since this can cause the sub-pixel localization routine to converge to the wrong position. For instance, a target on the SLS core can be partially or nearly occluded by the booster. Another common trigger is a target near the visual border between the model and the background.

3. Pose Determination

The final robustness improvement to the external calibration pipeline is the application of RANdom SAmple Consensus (RANSAC) [13] to the Perspective-n-Point (PnP) problem. PnP is the problem of estimating the relative pose (3D position and orientation) between a calibrated camera and an object given a set of n visible 3D points and their corresponding 2D projected image points [14]. In this case the object is the wind tunnel model, and the points are the targets. RANSAC ensures that even if there are incorrect matches between 3D targets and image locations, the final PnP solution will be valid. The algorithm is extremely robust and is used in many other computer vision and machine learning applications.

RANSAC is applied after the 3D targets have been matched to their image locations and is implemented by OpenCV [8] with the *solvePnPRANSAC* function. RANSAC is applied in that a subset of the matches are used to solve the PnP problem. If most of the other matches agree with the solution found using the subset, then that solution is kept. Agreement is defined by some threshold on reprojection error between the 3D target and image location.

B. **Data Collection Procedure**

1. 3D Target Position Measurement

The operational improvement suggested is to improve the process by which the (X, Y, Z) locations of the registration targets is acquired. PSP applications assume that the 3D positions of the targets are known and accurate since the positions of the targets are measured with a coordinate-measuring machine (CMM). However, the targets are still measured by a human operator attempting to manually locate the center of the target. If the measured location of a target is assumed to follow a zero-mean 2D normal distribution about the true center with a standard deviation of $1/12^{th}$ of its radius, the 3D target uncertainty would account for roughly 80% of the uncertainty significance. Targets in the most recent SLS wind tunnel tests had a radius of roughly 0.2” [2,4,5], which implies that with that assumed distribution the measured position is within 0.05” for 99.74% of targets. If anything, this assumption is overly generous since no indication is given for the center of the target. Additionally, this is assuming there is no additional error due to model deformation.

One mitigation to this issue is to have a center indication on the target masks. The target locations are masked with a circular vinyl patch during the paint application. These patches are typically laser cut since their diameter is custom. A very small circular cut, akin to a point, will be added to future patches to denote the center of the patch. The 3D position of the target will then be referenced to that point with the CMM. If this reduces the measurement uncertainty to 1/48th of a target radius (implying the measurement is within 0.0125” for 99.74% of targets), then the component of external calibration uncertainty from target position uncertainty is reduced by 75%.

C. **Processing Pipeline**

1. Keyframe Scheme

An ongoing endeavor is to implement a scheme conceptually similar to keyframes in video editing. Each process in the multiprocessing pipeline is given a set of several hundred sequential frames from a given test run. For this scheme, each process will select the center frame to be the keyframe and compute a two-stage external calibration for that frame. Every node on the model will then be checked for visibility, and the image locations of the visible nodes will be found using projection. For all other frames that the process is responsible for, a one-stage external calibration and the image locations are found using a gradient-based update.

For keyframes, a two-stage external calibration scheme was implemented. The initial guess using the commanded model position can be off by 5-8 pixels which makes detecting and correctly matching targets difficult. Only half of the visible targets are detected and matched, so the first stage is referred to as a coarse calibration. After the coarse calibration is complete, the improved model position is within 1 pixel, and nearly all targets can be correctly matched for the second stage of the calibration, the fine calibration. This fine calibration typically accurate to better than 0.05 pixels. Within a process, the nearby non-keyframes only need the fine calibration performed since model motion on the timescale of several hundred frames is small.

For keyframes, the full visibility checking process is performed, then projection with lens distortion for all visible nodes of the model mesh. This allows for all visible nodes to be associated with a pixel position. For nearby non-keyframes, it is assumed that the same nodes are visible. A gradient-based update using the Jacobian of the image position with respect to the translation and rotation vectors is used for non-keyframes instead of the projection process. This update is extremely fast, and the error introduced is much smaller than that from the sub-pixel localization, internal calibration, or 3D target positions.

2. Additional Target Types

In addition to the targets, the model also has Kulite pressure transducers (Kulites) that can be potentially used for external calibration. The masks used for the Kulites are smaller than the targets, measuring only 0.2” in diameter compared to 0.4” for targets. Since they are physically smaller, they occupy fewer pixels in the image and therefore have worse sub-pixel localization performance. The average Kulite localization uncertainty is 0.1 pixels compared to 0.06 pixels for the targets. However, there are roughly 60 Kulites visible to each camera (after removing Kulites too close together due to the robustness matching checks) compared to only 20 targets.

The external calibration uncertainty was quantified using Monte Carlo analysis using only the targets compared to the targets and Kulites. The expected external calibration uncertainty for only the targets was 0.1° and 0.22”. With the targets and Kulites, the uncertainty was reduced to 0.08° and 0.20”. The uncertainty was reduced, however there is a significant trade off. For one, the external calibration computation time increases when using both targets and Kulites. Additionally, the Monte Carlo analysis assumes that all matches are correct, and that the sub-pixel localization acts according to the quantified distribution. In practice, the sub-pixel localization routine often has trouble with the Kulites since surface artifacts can cause the localization routine to misfire. This can result in large errors that aren’t captured by the Monte Carlo analysis. One of these errors will cause the performance to be worse than using only the targets. Considering the uncertainty reduction is marginal, the drawbacks of using both targets and Kulites outweigh the benefits and only targets should be used.

IV. Future Work

Even with the suggested improvements to the external calibration pipeline, the uncertainty significance originating from the targets 3D position is a concern. Significant model deformation or planar targets such as with fundamental research experiments on a flat plate would lead to relatively high projection uncertainty. In future iterations of camera calibration software, it would be beneficial to implement object-releasing [15]. Object-releasing lifts the requirement of accurate 3D target positions and used multi-camera calibration to simultaneously estimate the external calibration and target geometry. This method would reduce the strict 3D target localization requirement and side-step the issue of estimating model deformation.

References

[1] Ruyten, W. *Real-Time Processing of Pressure-Sensitive Paint Images*. 2006.

[2] Roozeboom, N. H., Powell, J., Baerny, J. K., Murakami, D. D., Ngo, C. L., Garbeff, T. J., Ross, J. C., and Flach, R. "Development of Unsteady Pressure-Sensitive Paint Application on NASA Space Launch System." *AIAA Scitech 2019 Forum*, 2019, pp. 1–16. https://doi.org/10.2514/6.2019-3502.

[3] Sellers, M. E., Nelson, M. A., Roozeboom, N. H., and Burnside, N. J. "Evaluation of Unsteady Pressure Sensitive Paint Measurement Technique for Space Launch Vehicle Buffet Determination." *AIAA SciTech Forum*, No. January, 2017, pp. 1–34. https://doi.org/10.2514/6.2017-1402.

[4] Powell, J. M., Murman, S. M., Ngo, C. L., Roozeboom, N. H., Murakami, D. D., Baerny, J. K., and Li, J. "Development of Unsteady-PSP Data Processing and Analysis Tools for the NASA Ames Unitary 11ft Wind Tunnel." *AIAA Scitech 2020 Forum*, Vol. 1 PartF, No. January, 2020, pp. 1–17. https://doi.org/10.2514/6.2020-0292.

[5] Roozeboom, N., Baerny, J. K., Murakami, D. D., Ngo, C., and Powell, J. M. Recent Developments in NASA's Unsteady Pressure-Sensitive Paint Capability. 2020.

[6] Sellers, M. Advances in AEDC's Lifetime Pressure-Sensitive Paint Program. In *2005 US Air Force T\&E Days*, 2005, p. 7638.

[7] calib.io. Camera Calibration Boards, Targets and Software by Calib.Io. https://calib.io/.

[8] opencv dev team. Camera Calibration and 3D Reconstruction — OpenCV 2.4.13.7 Documentation. *OpenCV*. https://docs.opencv.org/2.4/modules/calib3d/doc/camera_calibration_and_3d_reconstruction.html.

[9] Ruyten, W., and Fisher, C. J. "Effects of Reflected Light in Luminescent Paint Measurements." *AIAA journal*, Vol. 39, No. 8, 2001, pp. 1587–1592.

[10] Le Sant, Y. Overview of the Self-Illumination Effect Applied to Pressure Sensitive Paint Applications. 2001.

[11] Meister, D., Ogaki, S., Benthin, C., Doyle, M. J., Guthe, M., and Bittner, J. A Survey on Bounding Volume Hierarchies for Ray Tracing. No. 40, 2021, pp. 683–712.

[12] Ruyten, W. "Model Attitude Determination in Wind Tunnel with a Luminescent Paint Data System." *AIAA journal*, Vol. 38, No. 9, 2000, pp. 1692–1697.

[13] Fischler, M. A., and Bolles, R. C. "Random Sample Consensus: A Paradigm for Model Fitting with Applications to Image Analysis and Automated Cartography." *Communications of the ACM*, Vol. 24, No. 6, 1981, pp. 381–395.

[14] NVIDIA Corporation. Isaac SDK — ISAAC 2020.1 Documentation. https://docs.nvidia.com/isaac/isaac/packages/pnp/doc/pnp.html.

[15] Strobl, K. H., and Hirzinger, G. More Accurate Pinhole Camera Calibration with Imperfect Planar Target. 2011.