

# Parametric Subpixel Matchpoint Recovery with Uncertainty Estimation: A Statistical Approach

Matt Steele and Christopher Jaynes \*  
Department of Computer Science  
University of Kentucky  
Lexington, KY, 40513

## Abstract

*We present a novel matchpoint acquisition method capable of producing accurate correspondences at subpixel precision. Given the known representation of the point to be matched, such as a projected fiducial in a structured light system, the method estimates the fiducial location and its expected uncertainty. Improved matchpoint precision has application in a number of calibration tasks, and uncertainty estimates can be used to significantly improve overall calibration results.*

*A simple parametric model captures the relationship between the known fiducial and its corresponding position, shape, and intensity on the image plane. For each matchpoint pair, these unknown model parameters are recovered using maximum likelihood estimation to determine a subpixel center for the fiducial. The uncertainty of the matchpoint center is estimated by performing forward error analysis on the expected image noise. Uncertainty estimates used in conjunction with the accurate matchpoints can improve calibration accuracy for multi-view systems.*

## 1. Introduction

Matchpoint estimation is a problem common to computer vision systems that exploit multiview geometry. Given multiple views of a scene, a sufficient number of corresponding image points can be used to recover canonical multi-view representations (i.e. the Fundamental matrix, and higher-order tensors). For multi-view systems, these geometric representations are the basis for a number of important computer vision tasks including structure and motion reconstruction, novel view synthesis, and photogrammetric analysis.

Regardless of the representation that the multi-view geometry will take, the pairs of matching points in each view must be computed. How this correspondence problem is solved ultimately influences the accuracy and robustness of

the estimated viewing geometry [23]. As a result, detecting corresponding points in images is the focus of an active research community. Approaches tend to be specific to the domain under consideration and vary significantly.

This work addresses the correspondence problem in active multi-view systems. Examples of these systems include structured light rigs [5, 7], multi-projector/camera displays [15], and dual passive and active laser-range finders [3]. In active systems, the difficulty of the correspondence problem is somewhat lessened due to the controllable nature of the light projectors involved [1]. However, accurate detection of emitted points of known shape is still an open problem.

We present a novel solution to this problem that produces accurate estimates of subpixel correspondence that surpass the accuracy of commonly used methods. A parametric model that describes the mapping between a fiducial and its corresponding position, shape, and intensity on the image plane is estimated. For each projected and observed matchpoint pair, these unknown parameters of the model are recovered using maximum likelihood estimation to determine a subpixel center for the fiducial. In addition, an estimate of the matchpoint uncertainty is produced.

Recent work has studied the statistical dependence of matchpoint accuracy on the accuracy of the calibration result [23, 17]. In particular, there is ample evidence that increasingly accurate matchpoints will result in an improved calibration result and that an estimate of matchpoint uncertainty can be taken into account during the calibration phase to further improve the result.

### 1.1. Related Work

Automatic estimation of point feature location and their corresponding positions in multi-view systems is a focus of research in both the computer vision and photogrammetric communities. In passive systems, interest operators [19] can produce potential matchpoints with known structure (i.e. significant gradient) and have been extended to produce subpixel estimates of point features in single views [10]. Similar approaches have been developed for

---

\*Work supported by NSF Research Infrastructure Award EIA-#0101242 and NSF CAREER Award EIA-#0092874

accurate estimation of targets of known shape, and can take into account whether its image will be affected by perspective distortion [18, 24]. Optical distortion can also be modeled as a way to improve calibration accuracy [13].

Automatic matching of these point features, or patterns of point features, involves the correlation of a particular image region with that of a corresponding region in a second view. The subpixel position of the matching point is typically given by the maximum of a function that has been fit to the local correlation values. Studies related to the accuracy of image correlation under these conditions have been conducted and are important in understanding how matchpoint accuracy influences epipolar geometry estimation and stereo reconstruction [23, 24, 12, 9]. Although this work focuses on generation of accurate matchpoints in active systems, the results have implications to ongoing efforts related to purely camera-based systems such as these.

Traditional active vision setups, such as structured-light systems, are metrically calibrated by placing a physical target in the field of view of all devices to provide a set of points whose world coordinates are precisely known [14, 7, 5]. More recent work has explored self-calibration methods that do not require point matches but these systems are forced to make significant assumptions [11], or are only capable of recovering a restricted number of degrees of freedom [6].

Relative calibration of camera-projector pairs is usually performed by projecting a target of known shape, (i.e. a Gaussian) that is subsequently detected in a second view [14, 3, 2]. The subpixel center of the matchpoint is recovered via least squares fitting to the observed intensities and selecting the fit maximum as the matchpoint center. Although this well-known method is capable of accurate matchpoint generation, it does not take into account perspective distortion and other properties of the system that influence matchpoint appearance. In contrast, the method presented here models the local distortion for each matchpoint candidate and consequently produces significantly more accurate matchpoint positions. In addition, the method produces an uncertainty estimate for each matchpoint that can be taken into account to further increase calibration accuracy [23].

In addition to calibration of structured-light scanners and other active systems, the work presented here has direct application to multi-projector display environments. Recently, a number of researchers have introduced camera-based calibration techniques that compute the relative or absolute calibration of clusters or projectors [4, 22, 21, 15]. Once calibrated, these systems can generate a uniform, large-format, potentially immersive image to a user. More accurate matchpoints as a result of the technique introduced here will improve the utility of these systems.




Figure 1: Perspective distortion and photometric characteristics of the camera-projector pair warp the appearance of projected fiducials. (a) Gaussian target in framebuffer of projector. (b) Same Gaussian projected on a planar surface and seen from a camera.

## 2. Matchpoint Recovery

For a given fiducial  $F$ , centered at projector pixel  $i, j$ , the role of matchpoint recovery is to detect the center of the projected fiducial in the camera to subpixel accuracy. Because  $F$  should be a fiducial with a low-autocorrelation coefficient for positional offsets and should change smoothly with respect to small offset, we fix  $F$  to be a circular Gaussian target centered at  $i, j$ .

Our approach to detecting the center of  $F$  in the frame of the camera is motivated by the observation that the appearance of  $F$  is warped due to perspective projection and sensor nonlinearities (See Figure 1).

As a result, we model the intensity of the projected target position, centered in the camera's image plane at  $x, y$  as:

$$I([x, y]) = C(F(\mathbf{H}[i, j, 1]^T)) \quad (1)$$

where  $C$  is a transfer function that maps projected color to observed colors,  $F$  is a function describing the template and  $\mathbf{H}$  is a geometric warp that maps points from the camera's image plane to the template's domain.

Note that  $\mathbf{H}$  governs the location of the template in the camera's image plane. Applying  $\mathbf{H}^{-1}$  to the template's origin yields the location of the template's origin in the camera. The essence of our approach is to estimate  $\mathbf{H}$  for a particular fiducial by finding the  $\mathbf{H}$  that maximizes the correlation between the right hand side of Equation 1 and the observed camera intensities,  $I([x, y])$ .

We assume that the surface is locally planar near the target, so  $\mathbf{H}$  is well approximated by a homography. Note that the projected target is circularly symmetric about its origin and any geometric warps which differ only by a rotation in the domain of the template (post-rotations) will produce identical fiducials in the frame of the camera. Consequently, we can eliminate one parameter from  $\mathbf{H}$ .

For simplicity, we model the color transfer function to be independent of pixel location or neighboring pixel values.

For the results presented here, we estimate the color transfer function a priori by projecting four different graylevels in sequence, and fitting the results to a four-parameter model. This method has been shown to accurately predict color/intensity values for common camera-projector setups [16].

Estimation of  $\mathbf{H}$ 's seven parameters is performed using an iterative optimization of the nonlinear function. An initial guess is determined by computing a bounding box for the pattern in the camera, and corresponding its corners with those of the unit square. In the following sections we present the technical details regarding the fitting process.

## 2.1 Technical Details

The color transfer function, which maps projector framebuffer intensities in the range  $[0, 255]$  to observed camera intensities in the same range, is parameterized in the following way:

$$C_{k_1, k_2, k_3, k_4}(I) = k_1 + \frac{k_2}{1 + e^{k_3 I + k_4}} \quad (2)$$

The four parameters  $k_1, k_2, k_3$  and  $k_4$  of the color transfer function  $C$  are estimated by projecting four different intensities and capturing the resulting image in a camera. These projected/observed pairs provide sufficient information to completely determine each parameter.

The template function  $F$ , which maps points in the plane to intensities in  $[0, 255]$ , is a circular Gaussian centered on the origin given by,

$$F([x, y]) = 255e^{-(x^2 + y^2)} \quad (3)$$

Because the Gaussian is centered at the origin, it must be translated and rescaled to yield an arbitrary target, centered at framebuffer pixel,  $x_{p_0}, y_{p_0}$ .

$$\vec{T}_{x_{p_0}, y_{p_0}, s}([x, y]) = \left[ \frac{x - x_{p_0}}{s}, \frac{y - y_{p_0}}{s} \right] \quad (4)$$

Because  $[x_{p_0}, y_{p_0}]$  will be the projector coordinates of the matchpoint, the higher level routine typically chooses these parameters. The parameter  $s$  determines the size of the target, and its value should be chosen to improve the ultimate accuracy of the results. Large values of  $s$  typically lead to longer processing times, and for non-flat display surfaces, can lead to accuracy problems, as the approximation of  $\mathbf{H}$  in Equation 1 as a homography becomes invalid. Other factors which argue for smaller  $s$  are radial distortion problems, and clutter resulting from a nonuniform display surface albedo. If  $s$  is too small, however, image noise should tend unduly affect the estimate of the camera's matchpoint. In Section 3 we demonstrate the accuracy of the system over several different scales of the projected Gaussian.

As noted in Section 1,  $\mathbf{H}$ , the mapping from the camera's frame to the target's domain is modeled as a homography.  $\mathbf{H}$  depends on the parameters for  $T$  chosen in Equation 4, on the relative positioning of the projector and camera, and on the location and orientation of the display surface.

$$\mathbf{H} = \begin{bmatrix} a & b & c \\ d & e & f \\ g & h & 1 \end{bmatrix} \quad (5)$$

$\mathbf{H}$  maps points into the domain of  $F$ , the circular Gaussian, which is invariant with respect to rotations about the origin. Consequently, a post rotation (or left multiply) can be applied to  $\mathbf{H}$  without affecting the correlation score. This results in a 3x3 matrix with a new parameterization, given by:

$$\begin{bmatrix} \cos \theta & \sin \theta & 0 \\ -\sin \theta & \cos \theta & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} a & b & c \\ d & e & f \\ g & h & 1 \end{bmatrix} = \begin{bmatrix} a \cos \theta + d \sin \theta & b \cos \theta + e \sin \theta & c \cos \theta + f \sin \theta \\ -a \sin \theta + d \cos \theta & -b \sin \theta + e \cos \theta & -c \sin \theta + f \cos \theta \\ g & h & 1 \end{bmatrix} \quad (6)$$

Upon inspection of equation 6, it should be obvious that if we let  $\theta = \arctan d/a$ , the term in the second row and first column vanishes. This implies that, if  $\mathbf{H}$  maximizes the correlation score, then there exists a matrix  $\tilde{\mathbf{H}}$ , whose second row and first column has a zero entry which also maximizes the correlation score.  $\tilde{\mathbf{H}}$  and  $\mathbf{H}$  produce identical camera matchpoints. This is true because the post-rotation takes place about  $[0, 0]$  in the template domain, system where the Gaussian is unwarped and centered on the origin.

$\mathbf{H}$  then is rewritten to exclude parameter  $d$  to yield a new matrix,  $\tilde{\mathbf{H}}$ . In addition, parameters  $c$  and  $f$  are written in terms of  $x_0$  and  $y_0$  so that the optimization of this new  $\tilde{\mathbf{H}}$  will allow us to solve for the uncertainty of the matchpoint center directly.

$$\tilde{\mathbf{H}} = \begin{bmatrix} a & b & -ax_0 - by_0 \\ 0 & e & -ey_0 \\ g & h & 1 \end{bmatrix} \quad (7)$$

Here, parameters  $c$  and  $f$  from Equation 5 are replaced by new parameters  $x_0$  and  $y_0$ .  $\tilde{\mathbf{H}}$  is defined so that  $\tilde{\mathbf{H}}[x_0 y_0 1]^T = [001]^T$ . That is,  $(x_0, y_0)$  are the image coordinates of the matchpoint center in the camera.

In practice, a background image from the camera is captured before each fiducial is projected. This is then subtracted from the scene containing the Gaussian target and a bounding box is fit to the largest remaining connected component. Optimization, then, only takes into account pixels




Figure 2: Bounding region fit to the Gaussian target. Initial estimates of the model parameters are derived from the four corners of the bounding region.

within the initial fit bounding box. This both improves performance and prevents data far from the matchpoint, which is more likely to be dominated by noise, from affecting the result. Figure 2 shows a closeup view of a projected fiducial and its automatically detected bounding box.

An initial estimate of the parameter vector  $[a, b, e, g, h, x_0, y_0]$ , is computed by corresponding the corners of the target’s bounding box in the camera’s image to the corners of the unit square in the domain of  $F$ . These four matchpoints determine an eight parameter homography,  $\hat{\mathbf{H}}$ .  $\hat{\mathbf{H}}$  can be computed by considering it as the composition of two warps, one which maps the corners of the target’s bounding box to the standard projective basis, and a second warp which maps the standard projective basis to the unit square. Each of these two warps can be computed by solving a 3x3 linear system [8].

In general, this initial guess,  $\hat{\mathbf{H}}$  has has eight parameters. In particular, entry  $d$  is potentially non-zero as required by our parameterization. Rather than applying a post-rotation to force  $d$  to be zero, we note that the homography corresponds to the corresponding points of two coordinate-axis aligned rectangles, so the resulting homography only consists of two scale parameters and two translation parameters, and entries  $b, d, g,$  and  $h$  will be zero.

Given an initial estimate of the warp,  $\hat{\mathbf{H}}$ , and the estimated color transfer function,  $C$ , the following sum of squared difference equation is minimized:

$$E(a, b, e, g, h, x_0, y_0) = \sum_{i,j} \left[ \frac{C(F(\hat{\mathbf{H}}[i, j, 1]^T)) - I(i, j)}{\sigma_{ij}} \right]^2 \quad (8)$$

where  $\sigma_{ij}$  is the variance for all pixel  $i, j$  in the bounding region under consideration. For the results shown here pixel variance was estimated by analyzing a sequence of images of a stationary scene.

To represent Equation 8 as a standard least-squares problem, we change our notation slightly, rewriting it as

$$E(a, b, e, g, h, x_0, y_0) = \sum_k \left[ \frac{y(\vec{x}_k; \vec{a}) - y_k}{\sigma_k} \right]^2 \quad (9)$$

Here,  $\vec{x}_k = [i, j, 1]^T$ , and  $\vec{a} = [a, b, e, g, h, x_0, y_0]^T$ , the optimization problem’s parameter vector. Note that in this case, the sum ranges over all  $k$  pixels contained in the bounding region to be fit. We drop the 2D,  $i, j$  indices for notational convenience.

We use an implementation of the Levenberg-Marquadt algorithm to minimize  $E$  of equation 8 [20]. Once convergence is reached, the camera’s matchpoint is given by the elements  $[x_0, y_0]$  of the parameter vector output by the optimizer.

## 2.2 Uncertainty Estimation

The optimization method, described in Section 2.1 converges at the  $\chi^2$  minimum of Equation 8. At that point the components of the Hessian matrix to Equation 9 is given by:

$$[\alpha] = \alpha_{kl} = \sum_{i=1}^N N \frac{1}{\sigma_i^2} \left[ \frac{\partial y(\vec{x}_i; \vec{a})}{\partial a_k} \frac{\partial y(\vec{x}_i; \vec{a})}{\partial a_l} \right], \forall k, l = 1..7. \quad (10)$$

The estimated covariance matrix  $[C]$ , of the errors of  $\mathbf{a}$  with respect to the fitted model is computed directly from the Hessian, given by,

$$[C] \equiv [\alpha]^{-1} \quad (11)$$

The diagonal elements of  $[C]$  contain the variances of the fit parameters  $\mathbf{a}$ . Because we have previously derived a form of the homography that explicitly denotes our subpixel center,  $x_0, y_0$  (see Equation 7). The covariances (or squared uncertainties) of these values are directly accessible without further analysis.

Uncertainty estimates that correspond to matchpoint centers, like those computed here, have been shown to improve overall performance of other numerical estimation methods encountered in multi-view geometry. Examples include more accurate estimation of the Fundamental matrix [23], and three-dimensional reconstruction. In future work, we hope to incorporate these particular error estimates into multi-view calibration methods to compute accurate epipolar geometry and higher order multi-view tensors.

## 3. Experimental Results

The technique was tested in simulation where noise, characteristics of the planar projective warp, and the color transfer functions could be controlled. Real world experiments,

involving a commodity light-projector and a digital video camera system were also conducted. In each case, we study the accuracy (and corresponding uncertainty) under various conditions. The method is compared to the common sub-pixel estimation approach in which an unwarped Gaussian is fit directly to the observed intensities.

### 3.1 Simulation

Simulated experiments were conducted in order to characterize the behavior of the method under controlled conditions. Given a target Gaussian  $F$ , it is first warped by applying  $\tilde{\mathbf{H}}$ , as defined by Equation 7 and a given set of warping parameters. The resulting pixels are color modified according to Equation 2 and given color transfer parameters. The algorithm was then applied to the resulting warped Gaussian to attempt to detect a subpixel center.

The system was tested by computing the matchpoint center for a given warped target Gaussian under differing levels of zero-mean Gaussian noise. Figure 3 depicts the accuracy, measured as the 2D distance from ground truth to estimated center. Noise levels range from zero standard deviations to 80.

The homography fit technique outperforms traditional Gaussian fitting for all levels of image noise. It is interesting to note that even at significant noise levels, the new approach produces matchpoint center estimates that are accurate to less than a pixel.

### 3.2 Real-world Tests

A number of experiments were conducted using a camera and projector pair with overlapping fields of view on a planar surface. Measuring the accuracy of matchpoint centers to subpixel accuracy under these conditions is a challenging task.

Baseline data for subpixel centers was acquired by projecting a large Gaussian target, with a standard deviation of 300 pixels, from the projector’s 1024x768 framebuffer. We also projected a very small target, with a standard deviation of one pixel, centered at the same position in the projector as the large Gaussian. Both images were taken with the projector camera pair in identical configurations. The small target, as seen from the camera, was used to manually estimate the center of both the large and small Gaussians. Subsampling the large Gaussian image and running the algorithm on several successively smaller resolution images produces a baseline subpixel offset for the centerpoint of the Gaussian in the downsampled image. Results of the algorithm can then be compared with the baseline estimates given by the subsampling process.

Each stage of subsampling halved the number of pixels in both the horizontal and vertical directions. We used the

following algorithm to perform the subsampling:

$$I_n(i, j) = \frac{1}{4} \left( \begin{array}{l} I_{n-1}(2i, 2j) + I_{n-1}(2i + 1, 2j) + \\ I_{n-1}(2i, 2j + 1) + I_{n-1}(2i + 1, 2j + 1) \end{array} \right) \quad (12)$$

where,  $I_n$  is the image after  $n$  stages of subsampling.

Given our subsampling pattern, the following recurrence relation describes pixel positions at scale  $n - 1$  to their corresponding position at scale  $n$ .

$$(i_n, j_n) = \left( \frac{i_{n-1}}{2} - \frac{1}{4}, \frac{j_{n-1}}{2} - \frac{1}{4} \right) \quad (13)$$

Iterating this recurrence leads to a series with the following closed-form solution.

$$(i_n, j_n) = \left( \frac{i_0}{2^n} - \frac{1 - \frac{1}{2^n}}{2}, \frac{j_0}{2^n} - \frac{1 - \frac{1}{2^n}}{2} \right) \quad (14)$$

At stage  $n$  of subsampling, we apply Equation 14 to the manual estimate of the matchpoint location on the full resolution image. We use this result as a baseline for evaluating our method of matchpoint location. The successive subsampling improves the quality of the baseline estimate; the uncertainty at stage  $n$  is  $1/2^n$  (pixels).


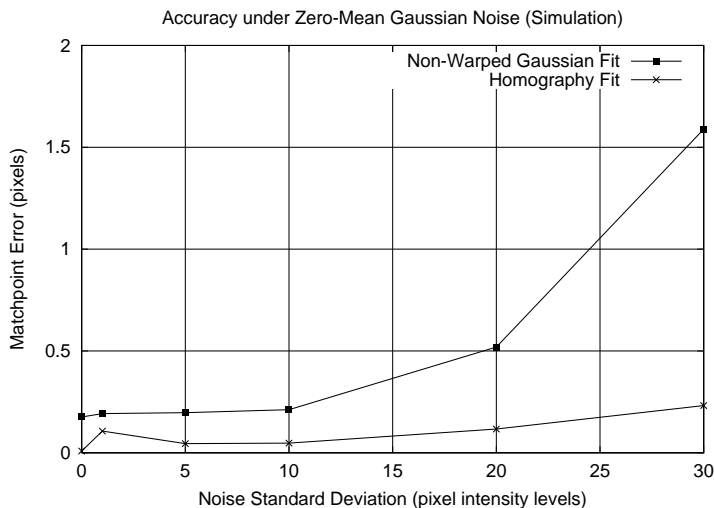
This method was used to analyze the accuracy method for several different relative projector-camera configurations. The different setups can be generally described as narrow and oblique. In the narrow case, the angle between the optic axis of each device varied between angles less than 45-degrees. Oblique setups refer to cases where the relative viewing angle exceeded 45-degrees.

For each setup, the method was applied to 10 different matchpoints and the mean distance between the baseline centerpoint and the estimated center was measured for all cases in each category. Table 1 reports the accuracy of the matchpoint center from fitting a Gaussian the the target without optimization versus the accuracy of the method introduced here.

Configuration	Gaussian Only	Homography Fit
Narrow	0.935	0.488
Oblique	0.877	0.398

Table 1: Error comparison for traditional matchpoint method versus technique presented in the paper. Errors reported are pixel distance from baseline to estimated target center. Results represent mean error computed over 10 trials.

The accuracy of the optimization method outperforms the traditional approach under each of the conditions tested. As the relative angle between the projector and camera is increased, the amount of warp on the fiducial is increased.



(a)

(b)

Figure 3: (a) Accuracy of the method versus increasing levels of image noise, as measured on simulated data. (b) Close-up of target corrupted by noise with a  $\sigma$  of 30.

As a result, the optimization technique that takes into account this potential change of fiducial appearance has the most impact. This is demonstrated by the Oblique case in Table 1.

Although the size of the initial bounding box increases optimization time, the accuracy of the method is stable with respect to the scale of the projected Gaussian. Figure 4 shows the accuracy of the method, for a particular camera-projector configuration, with respect to several different scales. In this case, the projector and camera were separated by approximately 30 degrees for moderate distortion of the Gaussian target.

Figure 4 also compares the optimizer’s expected uncertainties to the actual error obtained from the baseline data. Note that while the expected uncertainties appear to be useful, they consistently underestimate the error. While modeling the noise as zero-mean, uncorrelated and Gaussian appears to be somewhat reasonable, more accurate results and/or better uncertainty estimates might be obtained with a more complete noise model.

The most accurate results were typically obtained for the smallest template sizes, corresponding to a Gaussian target standard deviation of 16 pixels. At this scale, for the real-world tests, an accuracy of 0.3 pixels was reliably obtained.

## 4. Conclusion

This paper has introduced a matchpoint recovery method particularly designed for accurate subpixel estimation in active computer vision systems. The method pre-computes a color transfer function and optimizes seven parameters of a

warping homography that governs the appearance of a projected Gaussian target in the image plane. The parameters of the warping function are optimized in order to minimize the difference between the projected target and that seen in the camera. The matchpoint center is computed directly in this optimization phase.

Results demonstrate that the method produces more accurate matchpoint estimates than methods commonly employed by the computer vision community. In addition to improved accuracy, the technique produces a reasonable uncertainty estimate for each matchpoint pair. This information can be incorporated into traditional multi-view calibration methods.

In future work, we expect to explore more accurate initial estimates of the warping homography. Currently, initial estimates of this warp are supplied by a bounding box fit to the observed Gaussian. We hope to improve this estimate by fitting an ellipse to the warped Gaussian using level-set methods. In addition, we are exploring a more sophisticated model that simultaneously takes into account color transfer and local geometric warp. Finally, the method is being applied to passive sensing scenarios to accurately detect the center of known shapes that are subjected to similar locally planar warps. This approach promises to improve autocalibration of cameras observing a calibration target and more general vision tasks such as point feature tracking.

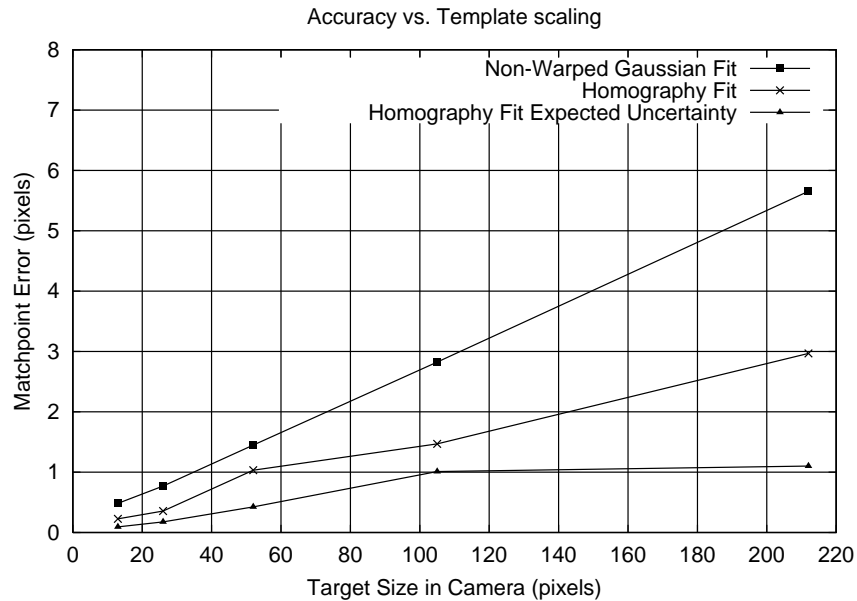


Figure 4: Accuracy of matchpoint center with respect to increasing scale. These results are all taken from the same oblique projector/camera configuration, involving significant foreshortening. Errors reported are the Euclidean distances from the baseline data.

## References

- [1] J. Battle, E. Mouaddib, and J. Salvi. Recent progress in coded structured light as a technique to solve the correspondence problem. *Pattern Recognition*, 31:963–982, 1998.
- [2] H. Chen, R. Sukthankar, G. Wallace, and T. Cham. Calibrating scalable multi-projector displays using camera homography trees. In *Computer Vision and Pattern Recognition*, 2001.
- [3] S. Chen and Y. F. Li. Self-recalibration of a colour-encoded light system for automated three-dimensional measurements. *Measuring Science and Technology*, 14:33–40, 2002.
- [4] Y. Chen, D. Clark, A. Finkelstein, T. Housel, and K. Li. Automatic alignment of high-resolution multi-projector displays using an un-calibrated camera. In *IEEE Visualization 2000*, pages 125–130, Salt Lake City, UT, October 2000.
- [5] Z. Chen and S.Y. Ho. Incremental model building of polyhedral objects using structured light. *Pattern Recognition*, 26/1:33–46, 1993.
- [6] C. Chu, S. Hwang, and Jung S. Calibration-free approach to 3d reconstruction using light stripe projections on a cube frame. In *IEEE 3rd Conference on 3D Digital Imaging and Modeling*, pages 13–19, 2001.
- [7] F. DePiero and M. Trivedi. 3d computer vision using structured light: design, calibration and implementation issues. *Advanced Computing*, 43:243–278, 1996.
- [8] O. D. Faugeras. *Three-Dimensional Computer Vision: A Geometric Approach*. MIT Press, 1993.
- [9] W. Foerstner. On the geometric precision of digital correlation. *International Archives of Photogrammetry and Remote Sensing*, XXIV(3):176–189, 1982.
- [10] W. Foerstner. A fast operator for detection and precise location of distinct points, corners and centers of circular features. In *ISPRS Intercommission Workshop on Fast Processing of Photogrammetric Data*, pages 281–305, 1987.
- [11] D. Fofi, J. Salvi, and E. Mouaddib. Uncalibrated vision based on structured light. In *IEEE Conference on Robotics and Automation*, 2001.
- [12] D. I. Havelock. The topology of locales and its effects on position uncertainty. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(4):380–386, 1991.

- [13] J. Heikkilae. Geometric camera calibration using circular control points. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(10):1066–1077, 2000.
- [14] D. Huynh. Calibration of a structured light system: a projective approach. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 225–230, 1997.
- [15] C. Jaynes, S. Webb, and R. M. Steele. A scalable framework for high-resolution immersive displays. *International Journal of the IETE*, 48(3 and 4):273–280, August 2002.
- [16] C. Jaynes, S. Webb, R. M. Steele, M. Brown, and B. Seales. Dynamic shadow removal from front projection displays. In *IEEE Visualization*, pages 175–182, San Diego, CA., October 2001.
- [17] J. Lin and J. Chuang. A geometry-based error estimation for cross-ratios. *Pattern Recognition*, 35(1):155–167, 2002.
- [18] E. M. Mikhail, M. Akey, and O. R. Mitchell. Detection and sub-pixel location of photogrammetric targets in digital images. *Photogrammetria*, 39(3):63–84, 1984.
- [19] H. Moravec. Rover visual obstacle avoidance. In *International Joint Conference on Artificial Intelligence*, pages 785–790, 1981.
- [20] W. Press, S. Teukolsky, W. Vetterling, and B. Flannery. *Numerical Recipes in C: Second Edition, Chapter 15: Modeling of Data*. Cambridge University Press, 1992.
- [21] R. Raskar, M. Brown, R. Yang, W. Chen, G. Welch, H. Towles, G. Seales, and H. Fuchs. Multi-projector displays using camera-based registration. In *IEEE Visualization*, pages 161–168, San Francisco, California, 1999.
- [22] R. Surati. *Scalable Self-Calibrating Display Technology for Seamless Large-Scale Displays*. PhD thesis, Computer Science and Electrical Engineering Department, Massachusetts Institute of Technology, 1999.
- [23] P. Torr and A. Zisserman. Performance characterization of fundamental matrix estimation under image degradation. *Machine Vision Applications*, 9(5-6):321–333, 1997.
- [24] J. C. Trinder, J. Jansa, and Huang Y. An assessment of the precision and accuracy of methods of digital target location. *Photogrammetria*, 50(2):12–20, 1995.