

Finding Corresponding Points Based On Bayesian Triangulation

Anand S. Bedekar and Robert M. Haralick
Dept. of Electrical Engineering, Box 352500,
University of Washington,
Seattle, WA 98195.
E-Mail:{anand,haralick}@ee.washington.edu

Abstract

In this paper, we consider the problems of finding corresponding points from multiple perspective projection images (the correspondence problem), and estimating the 3-D point from which these points have arisen (the triangulation problem). We pose the triangulation problem as that of finding the Bayesian maximum a posteriori estimate of the 3-D point, given its projections in N images, assuming a Gaussian error model for the image point co-ordinates and the camera parameters. We solve this by an iterative steepest descent method. We then consider the correspondence problem as a statistical hypothesis verification problem. Given a set of 2-D points, under the hypothesis that the points are in correspondence, the MAP estimate of the 3-D point is computed. Based on the MAP estimate, we derive a statistical test for verifying this hypothesis. To find sets of corresponding points when multiple points in each of N images are given, we propose a method that does the Bayesian triangulation and hypothesis verification on each N -tuple of points, selecting those that pass the hypothesis test. We characterize the performance of the Bayesian triangulation in terms of the average distance of the triangulated 3-D point from the true 3-D point, and of the point correspondence method in terms of its misdetection and false alarm rates.

Keywords: Bayesian, Correspondence Problem, Hypothesis Testing, Maximum A Posteriori Estimation, Triangulation.

1 Introduction

The point correspondence problem is the problem of finding sets of points which are projections of the same 3-D point, from multiple images of a scene. The triangulation problem is the problem of estimating the co-ordinates of a 3-D point, given its projections in multiple images. These problems arise in 3-D reconstruction, photogrammetry, object pose estimation, and point-feature-based object recognition.

In perspective projection imagery, a point in one image projects to an epipolar line on another image. If the camera parameters and the co-ordinates of the image points are known without error, then the correspondence problem can be solved by searching along the epipolar line of each point in the other images [2]. However, if there are random perturbations in the co-ordinates of the image points or the camera parame-

ters, no point will exactly be the intersection of the epipolar lines of its corresponding points. One could argue that if these perturbations are small, the intersections of the epipolar lines of corresponding points would be close to each other, and decide that a set of points is in correspondence if the separation of these intersections points is within a certain threshold. It is not clear how this threshold should be determined, however. In this sense the correspondence problem is not well-formulated, because the random perturbations in the camera parameters and the co-ordinates of the points have not been modelled explicitly.

The triangulation problem has been extensively analyzed in the photogrammetry literature [3][4]. If there is no error in the camera parameters and the image points, the 3-D point is the intersection in space of the rays projected through the image points from the respective centers of perspectivity. When there are random perturbations in the camera parameters and the image points, these rays will not intersect. In this case, the usual method is a least-squares approach, involving minimizing the sum of squared residual distances of the observed points from the projections of the triangulated point. In more recent work, Longuet-Higgins [5] derived a method for triangulation and estimating relative orientation given eight corresponding points from two images. Hartley [6] gave an analytical solution to the triangulation problem that is invariant to projective transformations of space, assuming a Gaussian noise model for the perturbations in the image points.

These triangulation methods make the assumption that the given 2-D points are guaranteed to be perspective projections of the same 3-D point. When the correspondences between the 2-D points are not known, one can make a hypothesis about a particular set of 2-D image points being in correspondence, and triangulate to get a 3-D point under this hypothesis. This hypothesis should then be validated, based on how good the estimate of the 3-D point is. However, the above methods do not address the issue of validating this hypothesis.

Recently, there has been considerable interest in the problem of finding corresponding image points with the aim of getting a good triangulation. Scott and Longuet-Higgins [7] proposed a method for finding corresponding points from a pair of images based on a proximity matrix that involves Gaussian weighted dis-

tances between features. They use the eigenvectors of this matrix to construct a pairing matrix, which minimizes the inner product of the pairing matrix and the proximity matrix. Shapiro and Brady [8] also proposed a similar eigenvector approach to determining point correspondences. Cheng *et al* [9] proposed two methods that find corresponding points and do triangulation. One approach is based on the work of Scott and Longuet-Higgins [7]. They construct a proximity matrix using the squared distances between the observed 2-D points and the projections of a "pseudo-triangulated" point, that is, a point that minimizes the 3-D distance to the rays projected back through the image points. The other method poses the correspondence problem as that of finding a maximum matching in a bipartite graph in which the nodes are the points from two images. This problem is then solved as a network flow maximization problem by assigning an appropriate weight to each edge based on a similarity function of the points that are connected by the edge.

None of these methods takes into account the uncertainty in the locations of the image points and the camera parameters with an explicit noise model. Also these methods are restricted to finding corresponding points and triangulating them from just two images, and do not generalize to more than two images.

In this paper, we find the maximum *a posteriori* estimate of the 3-D point, given its projections in N images, assuming a Gaussian error model for the image point co-ordinates and the camera parameters. We then consider the correspondence problem as a statistical hypothesis verification problem. Given a set of 2-D points, under the hypothesis that the points are in correspondence, the MAP estimate of the 3-D point is computed. Based on this estimate, we derive a statistical test for verifying the hypothesis that the given image points are in correspondence. To find sets of corresponding points when multiple points in each of N images are given, we propose a method that does the Bayesian triangulation and hypothesis verification on each N -tuple of points, and selects those that pass the hypothesis test. We also describe extensive experiments that were performed to validate the theory, and to characterize the performance of the point correspondence algorithm.

2 A Bayesian Formulation of the Triangulation Problem

In this section, we formulate the triangulation problem as a Bayesian maximum *a posteriori* estimation problem, assuming the perturbations in the observed 2-D points and the camera parameters to be Gaussian distributed with known covariance matrices.

2.1 Problem Statement

Let $\hat{x}_1, \hat{x}_2, \dots, \hat{x}_N$ be observed image points, from perspective projection images I_1, \dots, I_N respectively. Estimates $\hat{\theta}_1, \hat{\theta}_2, \dots, \hat{\theta}_N$ of the parameter vectors of the cameras of images I_1, \dots, I_N respectively are given. We assume that $\hat{\theta}_1, \hat{\theta}_2, \dots, \hat{\theta}_N$ are independent and normally distributed around the true camera parameters $\theta_1, \dots, \theta_N$ respectively, with known covariance matrices $\Sigma_{\hat{\theta}_1}, \dots, \Sigma_{\hat{\theta}_N}$ respectively. We

consider the true camera parameters $\theta_1, \dots, \theta_N$ as random variables with independent *a priori* densities $p(\theta_1), \dots, p(\theta_N)$ respectively. The observed points are the result of random perturbations on the perspective projections of a 3-D point \tilde{q} in the respective images, i.e.

$$\hat{x}_i = P(\tilde{q}, \theta_i) + \xi_i, i = 1, \dots, N$$

where $P(\tilde{q}, \theta_i)$ denotes the perspective projection of the 3-D point \tilde{q} in a camera with parameter vector θ_i . We assume the random perturbations $\xi_i, i = 1, \dots, N$ to be independent, and Gaussian distributed with zero mean and known covariance matrices $\Sigma_{\hat{x}_1}, \dots, \Sigma_{\hat{x}_N}$ respectively. We consider the 3-D point \tilde{q} as a random variable with a *a priori* density $p(\tilde{q})$. The triangulation problem can then be posed as:

Find $q = (x, y, z)^T$ to maximize

$$p(q | \hat{x}_1, \dots, \hat{x}_N, \hat{\theta}_1, \dots, \hat{\theta}_N)$$

which is equivalent to finding $q = (x, y, z)^T$ to minimize

$$-\ln p(q, \hat{x}_1, \dots, \hat{x}_N, \hat{\theta}_1, \dots, \hat{\theta}_N)$$

2.2 Solution

We have,

$$\begin{aligned} p(q, \hat{x}_1, \dots, \hat{x}_N, \hat{\theta}_1, \dots, \hat{\theta}_N) \\ = \int_{\theta_1} \dots \int_{\theta_N} p(\hat{x}_1, \dots, \hat{x}_N | q, \hat{\theta}_1, \dots, \hat{\theta}_N, \theta_1, \dots, \theta_N) \cdot \\ p(q, \hat{\theta}_1, \dots, \hat{\theta}_N | \theta_1, \dots, \theta_N) p(\theta_1, \dots, \theta_N) d\theta_1 \dots d\theta_N \end{aligned}$$

Given the true camera parameters $\theta_1, \dots, \theta_N$, q is independent of $\hat{\theta}_1, \dots, \hat{\theta}_N$, hence

$$\begin{aligned} p(q, \hat{x}_1, \dots, \hat{x}_N, \hat{\theta}_1, \dots, \hat{\theta}_N) \\ = p(q) \prod_{i=1}^N \int_{\theta_i} p(\hat{x}_i | q, \theta_i) p(\hat{\theta}_i | \theta_i) p(\theta_i) d\theta_i \end{aligned}$$

Here we also used the facts that the perturbations on the points $\hat{x}_1, \dots, \hat{x}_N$ are independent of each other, the perturbations in $\hat{\theta}_1, \dots, \hat{\theta}_N$ are independent of each other, and $\theta_1, \dots, \theta_N$ are *a priori* independent of each other.

Now, if the covariance matrices $\Sigma_{\theta_1}, \dots, \Sigma_{\theta_N}$ have small diagonal entries (as is usually the case), then to a first approximation,

$$\int_{\theta_i} p(\hat{x}_i | q, \theta_i) p(\hat{\theta}_i | \theta_i) p(\theta_i) d\theta_i \approx p(\hat{x}_i | q, \hat{\theta}_i) p(\hat{\theta}_i)$$

Hence

$$p(q, \hat{x}_1, \dots, \hat{x}_N, \hat{\theta}_1, \dots, \hat{\theta}_N) \approx p(q) \prod_{i=1}^N p(\hat{x}_i | q, \hat{\theta}_i) p(\hat{\theta}_i)$$

To find $p(\hat{\mathbf{x}}_i | \mathbf{q}, \hat{\theta}_i)$, we note that

$$\begin{aligned} \hat{\mathbf{x}}_i - P(\mathbf{q}, \hat{\theta}_i) &= \hat{\mathbf{x}}_i - P(\mathbf{q}, \theta_i) + P(\mathbf{q}, \theta_i) - P(\mathbf{q}, \hat{\theta}_i) \\ &\approx \hat{\mathbf{x}}_i - P(\mathbf{q}, \theta_i) + \left[\frac{\partial P}{\partial \theta}(\mathbf{q}, \hat{\theta}_i) \right] (\theta_i - \hat{\theta}_i) \end{aligned} \quad (1)$$

By our noise model, $\hat{\mathbf{x}}_i - P(\mathbf{q}, \theta_i) \sim N(\mathbf{0}, \Sigma_{\hat{\mathbf{x}}_i})$, and $(\theta_i - \hat{\theta}_i) \sim N(\mathbf{0}, \Sigma_{\hat{\theta}_i})$. Hence $\hat{\mathbf{x}}_i - P(\mathbf{q}, \hat{\theta}_i) \sim N(\mathbf{0}, \hat{\Sigma}_i(\mathbf{q}, \hat{\theta}_i))$, where

$$\hat{\Sigma}_i(\mathbf{q}, \hat{\theta}_i) = \Sigma_{\hat{\mathbf{x}}_i} + \left[\frac{\partial P}{\partial \theta}(\mathbf{q}, \hat{\theta}_i) \right] \Sigma_{\hat{\theta}_i} \left[\frac{\partial P}{\partial \theta}(\mathbf{q}, \hat{\theta}_i) \right]^T$$

Hence

$$p(\hat{\mathbf{x}}_i | \mathbf{q}, \hat{\theta}_i) = N \left(P(\mathbf{q}, \hat{\theta}_i), \hat{\Sigma}_i(\mathbf{q}, \hat{\theta}_i) \right)$$

So finding \mathbf{q} to minimize

$$-\ln p(\mathbf{q}, \hat{\mathbf{x}}_1, \dots, \hat{\mathbf{x}}_N, \hat{\theta}_1, \dots, \hat{\theta}_N)$$

is equivalent to finding \mathbf{q} to minimize

$$\begin{aligned} &-\ln p(\mathbf{q}) \\ &+ \sum_{i=1}^N \left\{ \frac{1}{2} (\hat{\mathbf{x}}_i - P(\mathbf{q}, \hat{\theta}_i))^T [\hat{\Sigma}_i(\mathbf{q}, \hat{\theta}_i)]^{-1} (\hat{\mathbf{x}}_i - P(\mathbf{q}, \hat{\theta}_i)) \right. \\ &\left. + \frac{1}{2} \ln 2\pi + \frac{1}{2} \ln |\hat{\Sigma}_i(\mathbf{q}, \hat{\theta}_i)| - \ln p(\hat{\theta}_i) \right\} \end{aligned}$$

If the 3-D point and the camera parameters are *a priori* uniformly distributed, $p(\mathbf{q})$ and $p(\hat{\theta}_i), i = 1, \dots, N$ are constants. Hence the problem reduces to finding \mathbf{q} to minimize

$$\begin{aligned} \epsilon(\mathbf{q}, \hat{\mathbf{x}}_1, \dots, \hat{\mathbf{x}}_N, \hat{\theta}_1, \dots, \hat{\theta}_N) &= \sum_{i=1}^N (\hat{\mathbf{x}}_i - P(\mathbf{q}, \hat{\theta}_i))^T [\hat{\Sigma}_i(\mathbf{q}, \hat{\theta}_i)]^{-1} (\hat{\mathbf{x}}_i - P(\mathbf{q}, \hat{\theta}_i)) \\ &+ \ln |\hat{\Sigma}_i(\mathbf{q}, \hat{\theta}_i)| \end{aligned} \quad (2)$$

We do the minimization by a steepest descent procedure. We start with an initial guess $\mathbf{q}_0 = (x_0, y_0, z_0)^T$, and at each iteration i , take a step $(\Delta x_{i+1}, \Delta y_{i+1}, \Delta z_{i+1})^T$ from \mathbf{q}_i in the direction of the negative gradient of ϵ evaluated at \mathbf{q}_i , so as to reach the minimum of ϵ along that direction. We use a golden search [10] to find the minimum in the negative gradient direction. The intuitive idea behind this is that close to the minimum, $\hat{\Sigma}_i(\mathbf{q}, \hat{\theta}_i)$, and hence $\ln |\hat{\Sigma}_i(\mathbf{q}, \hat{\theta}_i)|$ term in ϵ , is almost a constant, and so the variation in ϵ is dominated by the variation in

$$\sum_{i=1}^N (\hat{\mathbf{x}}_i - P(\mathbf{q}, \hat{\theta}_i))^T [\hat{\Sigma}_i(\mathbf{q}, \hat{\theta}_i)]^{-1} (\hat{\mathbf{x}}_i - P(\mathbf{q}, \hat{\theta}_i))$$

This sum of quadratic forms should be a convex \cup function, and has a unique minimum in the negative gradient direction.

In the implementation, the vector of camera parameters was taken as

$$\theta = (a, b, c, x_0, y_0, z_0, f)^T$$

(a, b, c) are the quaternion parameters [11] of the rotation matrix that rotates the world reference frame into the camera reference frame. (x_0, y_0, z_0) is the origin of the camera reference frame in the world co-ordinate system. f is the focal length of the camera. Details about the computation of the perspective projection $P(\mathbf{q}, \theta)$ of a 3-D point \mathbf{q} , the gradient $\frac{\partial \epsilon}{\partial \mathbf{q}}$ of the objective function of the Bayesian triangulation, and the computation of an initial guess for the steepest descent solution, using this representation, are given in [1].

3 Finding Corresponding Points

Suppose N points are given, one in each of N perspective projection images, and the question is to determine whether or not the points are in correspondence. To answer this, one can make a hypothesis that the points are in correspondence, and triangulate these points to obtain an estimate of the 3-D point. This estimate should then be used to verify the hypothesis that the given 2-D points were in correspondence. The Bayesian formulation of the triangulation problem described in the previous sections leads to such a test for verifying the hypothesis that the given 2-D points are in correspondence, as follows.

Because of the first order approximation we made in equation (1), the MAP estimate of \mathbf{q} under a uniform prior density is equivalent to the maximum likelihood estimate under the following noise model:

$$\hat{\mathbf{x}}_i = P(\mathbf{q}, \hat{\theta}_i) + \xi_i, i = 1, \dots, N$$

where $\xi_i \sim N(\mathbf{0}, \hat{\Sigma}_i(\mathbf{q}, \hat{\theta}_i))$. Close to the minimum, this nonlinear model can be approximated to a first order by a linear model. Consider a linear model of the form

$$\mathbf{y} = A\mathbf{x} + \eta$$

where \mathbf{y} is an M -dimensional vector of observations, \mathbf{x} is a p -dimensional parameter vector, and $\eta \sim N(\mathbf{0}, \Sigma)$. Let $\hat{\mathbf{x}}$ be the maximum likelihood estimate of \mathbf{x} . It is known [12] that the quadratic form

$$(\mathbf{y} - A\hat{\mathbf{x}})^T \Sigma^{-1} (\mathbf{y} - A\hat{\mathbf{x}})$$

has a χ_{M-p}^2 distribution. This quadratic form is used as a statistic to verify the hypothesis that the model is correct. So if the hypothesis that the 2-D points are in correspondence were correct, then at the point \mathbf{q} that minimizes ϵ ,

$$\delta(\mathbf{q}, \hat{\mathbf{x}}_1, \dots, \hat{\mathbf{x}}_N, \hat{\theta}_1, \dots, \hat{\theta}_N) = \sum_{i=1}^N (\hat{\mathbf{x}}_i - P(\mathbf{q}, \hat{\theta}_i))^T [\hat{\Sigma}_i(\mathbf{q}, \hat{\theta}_i)]^{-1} (\hat{\mathbf{x}}_i - P(\mathbf{q}, \hat{\theta}_i)) \quad (3)$$

would have a χ_{2N-3}^2 distribution. This is because $\hat{x}_1, \dots, \hat{x}_N$ are all 2-D vectors, and the vector \mathbf{q} being estimated is a 3-D vector. $\delta(\mathbf{q}, \hat{x}_1, \dots, \hat{x}_N, \hat{\theta}_1, \dots, \hat{\theta}_N)$ can thus be used as a test statistic to verify the hypothesis that the given 2-D points are in correspondence. Let x_α be such that $P(\chi_{2N-3}^2 > x_\alpha) = \alpha$. If $\delta > x_\alpha$, reject the hypothesis that the N points are in correspondence. The probability of rejecting a true hypothesis in this statistical test is α .

To find corresponding N -tuples of points when multiple points in each image are given, we propose the following method. For each N -tuple of points containing one point from each image, under the hypothesis that those N points are in correspondence, we do the Bayesian triangulation and hypothesis verification. If an N -tuple passes the test, we declare that set of points as a correspondence. The threshold for the hypothesis test can be set to limit the probability of rejecting a correct correspondence to a given value.

4 Experiments and Results

In this section we describe experiments to characterize the performance of the Bayesian triangulation in terms of the distance of the triangulated 3-D point to the true 3-D point, and of the point correspondence method in terms of its misdetection and false alarm rates. In these experiments, we used the geometry and camera parameters of model board 2 of the RADIUS data set [13]. The parameters used for the cameras and their covariance matrices of the were estimated by a multi-image camera calibration procedure [14]. For characterizing the performance of the Bayesian triangulation, 3-D points were randomly generated from the volume occupied by the model board, and projected into the cameras. The co-ordinates of the projected points and the camera parameters were perturbed with Gaussian noise of the respective covariance matrices. The noisy points and camera parameters were then input to the Bayesian triangulation and hypothesis test procedure.

Fig. 1 plots the mean and the standard deviation of the distance of the triangulated 3-D point from the true 3-D point as a function of the variance of the noise in the image points. Fig. 2 plots the mean and the standard deviation of the distance of the triangulated 3-D point from the true 3-D point, as a function of the number of cameras used for the triangulation. 10000 trials were done for each value of the noise variance and the number of cameras. For model board 2 of the RADIUS data set, the cameras are approximately 20 feet from the board. It is seen that the average distance from the estimated 3D point to the true 3D point is about 0.4 inches using 3 cameras, and reduces to about 0.1 inch using 15 cameras for the triangulation. Increasing the noise does not significantly degrade this performance.

The misdetection rate of the hypothesis test with a given threshold is the ratio of the number of trials for which the value of the test statistic is more than the threshold to the total number of trials. Fig. 3 plots the misdetection rate as a function of the threshold of the test, for various number of cameras used for the tri-

Mean and Std.Dev. of Distance of Triangulated 3-D Point from True 3-D Point




Figure 1: Mean and standard deviation of the distance of the triangulated 3-D point from true 3-D point v/s noise variance.

Mean and Std.Dev. of Distance of Triangulated 3-D Point from True 3-D Point




Figure 2: Mean and standard deviation of the distance of the triangulated 3-D point from true 3-D point v/s number of cameras.

angulation. The threshold is the corresponding quantile of the χ_{2N-3}^2 distribution. Thus the misdetection rate is seen to be 0.2 for a threshold of 0.8, 0.3 for a threshold of 0.7, etc. This also validates the assertion that the distribution of the statistic is χ_{2N-3}^2 . This distribution was also validated from the histograms of the test statistic by the Kolmogorov-Smirnov test [10]. Fig. 4 plots the misdetection rate as a function of the variance of the noise in the image points. The misdetection rate is seen to be almost invariant to the noise variance. This is because the distribution of the test statistic is χ_{2N-3} , independent of the noise variance.

We characterized the performance of the point correspondence method proposed in section 3 in terms of its false alarm and misdetection rates. An N -tuple of points that is declared to be in correspondence by the point correspondence method, but is actually not in correspondence, is a false alarm. The ratio of the




Figure 3: Misdetecation rate of hypothesis test v/s threshold.

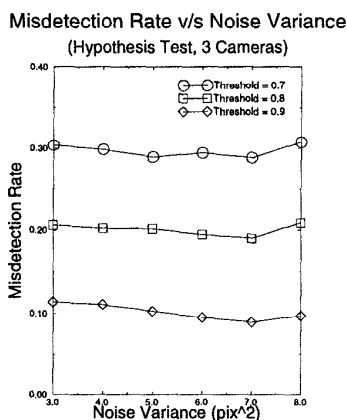


Figure 4: Misdetecation rate of hypothesis test v/s noise variance.

number of false alarms produced by the algorithm to the total number of incorrect correspondences input to the algorithm is the false alarm rate. An N -tuple of points which are actually projections of the same 3-D point, but which is not declared as a correspondence by the point correspondence method, is a misdetection. The ratio of the number of misdetections produced by the algorithm to the total number of true correspondences input to the algorithm is the misdetection rate. In these experiments, 20 3-D points were randomly generated from the volume of model board 2 of the RADIUS data set, and projected in 3 cameras. The projected points and camera parameters were perturbed by Gaussian noise. The lists of points in each image were randomly shuffled and input to the point correspondence method, which does the Bayesian triangulation and hypothesis verification on each triple of points, accepting those that pass the test as valid correspondences. 10 such trials were done for each value of the variance of the noise in the image points.

Fig. 5 plots the misdetection rate of the point corre-

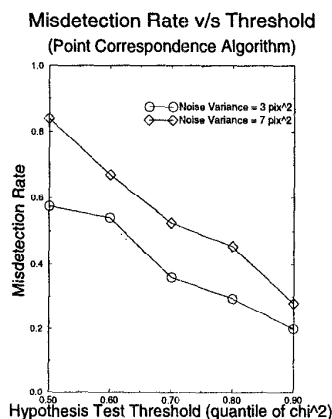


Figure 5: Misdetecation rate of point correspondence method v/s threshold.

spondence method as a function of the threshold used in the hypothesis test, and Fig. 6 plots the false alarm

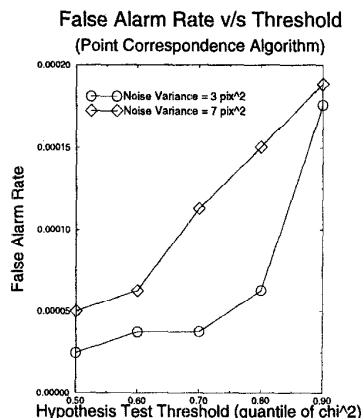


Figure 6: False alarm rate of point correspondence method v/s threshold.

rate against the threshold. The misdetection rate was somewhat higher than predicted by the threshold of the hypothesis test. A possible reason for this is that the number of trials (10 different noise values for the perturbations in the points and the camera parameters) was not sufficient to give experimentally observed misdetection rates close to the predicted rate. Fig. 7 plots misdetection rate v/s false alarm rate operating curves.

5 Discussion and Conclusion

In this paper, we gave a statistically sound formulation of the problems of deciding whether a given set of N points from N perspective projection images is in correspondence, and triangulating these points to estimate the 3-D point that gave rise to them. We posed

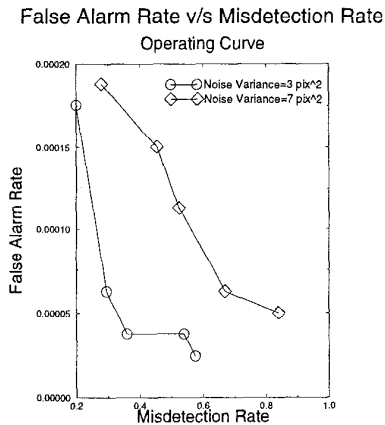


Figure 7: Misdetection rate v/s false alarm rate.

the triangulation problem as a Bayesian maximum *a posteriori* estimation problem, assuming a Gaussian error model for the observed 2-D points and the camera parameters. We solved this problem by an iterative steepest descent method. Based on this Bayesian estimate of the triangulated 3-D point, we derived a statistical test for verifying the hypothesis that the given 2-D points were in correspondence. The threshold of this hypothesis test can be set to limit the misdetection rate to a desired level. For finding corresponding points when multiple points in each of N images are given, we proposed a method that does this hypothesis verification test with a given misdetection rate threshold on each N -tuple of points having one point from each image, and accepts those N -tuples that pass the test as valid correspondences. We carried out extensive experiments to characterize the performance of the Bayesian triangulation and the point correspondence method.

A drawback of this method of finding corresponding points is that it is computationally very slow. To reduce the number of tuples of points that have to be examined, a fast front-end to this algorithm should be developed for generating reasonable hypotheses. This can be done, for instance, by projecting each image point to its epipolar lines on the other images, and collecting points which lie "close" to each others' epipolar lines. The Bayesian triangulation and hypothesis verification will then be done only on these hypothesized correspondences. This front-end could also be any other method of finding corresponding points. The correspondences generated by this method can then be considered as hypotheses verified by the hypothesis test.

The Bayesian triangulation and the point correspondence method are at present being used in conjunction with camera calibration [14], feature extraction operators [15][16], and partial model reconstruction algorithms [17] in the RADIUS project for automatic site model construction.

Acknowledgements

The authors would like to thank Ken Thornton and Xufei Liu for many helpful discussions and suggestions.

References

- [1] Anand Bedekar, "Finding corresponding points based on Bayesian triangulation", Master's thesis, Department of Electrical Engineering, University of Washington, 1995.
- [2] R.M.Haralick and L.G.Shapiro, *Computer and Robot Vision*, vol. II, Reading, MA:Addison-Wesley, 1992.
- [3] G.H.Schut, "An analysis of methods and results in analytical aerial triangulation", *Photogrammetria*, , no. 1, pp. 16-32, 1957.
- [4] H. Schmid, "An analytical treatment of the problem of triangulation by stereophotogrammetry", *Photogrammetria*, , no. 3, pp. 66-77 and 91-116, 1956.
- [5] H.C.Longuet-Higgins, "A computer algorithm for reconstructing a scene from two projections", *Nature*, vol. 293, pp. 133-135, Sept. 1981.
- [6] R.I.Hartley and P.Sturm, "Triangulation", in *Proc. ARPA Image Understanding Workshop*, 1994, pp. 957-966.
- [7] G.L.Scott and H.C.Longuet-Higgins, "An algorithm for associating the features of two patterns", in *Proc. Royal Society, London*, 1991, vol. B244, pp. 21-26.
- [8] L.S.Shapiro and J.M.Brady, "Feature-based correspondence: an eigenvector approach", *Image and Vision Computing*, vol. 10, no. 5, pp. 283-288, June 1992.
- [9] Y-Q.Cheng, R.T.Collins, A.R.Hanson, and E.M.Riseman, "Triangulation without correspondences", in *Proc. ARPA Image Understanding Workshop*, 1994, pp. 993-1000.
- [10] W.H.Press *et al*, *Numerical Recipes in C: The Art of Scientific Computing*, Cambridge University Press, 2nd edition, 1992.
- [11] R.M.Haralick and L.G.Shapiro, *Computer and Robot Vision*, vol. II, Reading, MA:Addison-Wesley, 1992.
- [12] C.E.Lunneborg and R.D.Abbott, *Elementary multivariate analysis for the behavioral sciences : applications of basic structure*, New York:North-Holland, 1983.
- [13] Ken Thornton *et al*, "Groundtruthing the RADIUS model-board imagery", in *Proc. ARPA Image Understanding Workshop*, 1994.
- [14] Kenneth B. Thornton and Robert M. Haralick, "Multi-image camera calibration and object-point estimation", in *Integrating Photogrammetric Techniques with Scene analysis and Machine Vision II*. SPIE, 1995, vol. 2486.
- [15] V.Ramesh and R.M.Haralick, "An integrated gradient edge detector - theory and performance evaluation", in *Proc. ARPA Image Understanding Workshop*, 1994, pp. 689-695.
- [16] X.Zhang, V.Ramesh, R.M.Haralick, and Anand Bedekar, "A Bayesian corner detector: Theory and performance evaluation", in *Proc. ARPA Image Understanding Workshop*, 1994, pp. 703-715.
- [17] X.Liu, R.M.Haralick, and Anand S. Bedekar, "Optimization methods for estimating 3-D object parameters", in *Proc. ARPA Image Understanding Workshop*, 1994, pp. 1017-1027.
- [18] Anand S.Bedekar and Robert M.Haralick, "A Bayesian method for triangulation and its application to finding corresponding points", in *Proc. IEEE Intl Conf. Image Processing*, 1995, vol. II, pp. 362-365.