

Shape recognition from three-dimensional point measurements with range and direction uncertainty

Xin Zhou

Michael D. DeVore, MEMBER SPIE

University of Virginia

Systems and Information Engineering

P.O. Box 400747

151 Engineer's Way

Charlottesville, Virginia 22904

E-mail: mdevore@virginia.edu

Abstract. We derive a pair of algorithms, one optimal and the other approximate, for recognizing three-dimensional objects from a collection of points chosen from their surface according to some probabilistic mechanism. The measurements are assumed to be noisy, and the measured location of a given point is translated according to a noise probability distribution. Distributions governing surface point selection and measurement noise can take a variety of forms depending upon the particular measurement scenario. At one extreme, each measurement is assumed to yield values restricted to a one-dimensional ray, a special case commonly adopted in the literature. At the other extreme, measured points are chosen uniformly from the object's surface, and the noise distribution is spherically symmetric, a worst-case scenario that involves no prior information about the measurements. We apply these two algorithms to shape recognition problems involving simple geometrical objects, and examine their relative behavior using a combination of analytical derivation and Monte Carlo simulation. We show that the approximate algorithm can be far simpler to compute, and its performance is competitive with the optimal algorithm when noise levels are relatively low. We show the existence of a critical noise level, beyond which the approximate algorithm exhibits catastrophic failure. © 2005 Society of Photo-Optical Instrumentation Engineers. [DOI: 10.1117/1.2138067]

Subject terms: pattern recognition; laser radar; machine vision; computer vision.

Paper 040792R received Oct. 21, 2004; revised manuscript received May 4, 2005; accepted for publication May 16, 2005; published online Dec. 8, 2005.

1 Introduction

The focus of this paper is on recognizing objects in a scene on the basis of the measured locations of randomly chosen points on the surface of an object. A wide variety of systems have been developed over the last several decades that measure the three-dimensional (3-D) location of points on the surface of objects. These include stylus-based contact methods¹ as well as noncontact methods such as laser radar (ladar), interferometric synthetic aperture radar,² and photometric techniques (for example, shape-from-motion,³ structured light methods,⁴ or a combination⁵). Traditional applications of 3-D data collections include building models of prototype objects for computer-aided design and manufacturing (CAD/CAM), rendering 3-D scenes in computer graphics, and collision avoidance for mobile robots. For a historical survey of data collection, object representation, techniques, and applications see Besl and Jain.⁶

Recently, considerable attention has focused on exploiting the rich information content in 3-D measurements for use in object recognition. Much of this has been motivated by the increasing application of ladar sensing platforms. For example, Zheng et al.⁷ compare range measurements on a pixel-by-pixel basis with values that follow from a hypothesized model and assign an overall score based on the degree to which object silhouette and surface features

match. Hutchinson et al.⁸ use a similar match score in a series of simulation experiments varying range and cross-range resolution and for different degrees of occlusion, with range measurement error simulated by additive Gaussian noise.

Development of an explicit statistical model for point measurement systems makes possible likelihood-based approaches. Green and Shapiro⁹ develop an object detection algorithm based on a statistical model for ladar range measurement error that accounts for range errors due to both limited sensor resolution and measurement anomalies arising from speckle fading. It is assumed that the direction vectors corresponding to the line of sight are known for each pixel, so the positional uncertainty for each measured point is restricted to a one-dimensional subspace. Lanterman¹⁰ builds on this model by allowing local range accuracy and probability of measurement anomalies to be functions of range. Further, his model accommodates wrapping of the measured range across range ambiguity intervals as can happen with coherent heterodyne-detection systems. Whitaker and Gregor¹¹ use a truncated distribution based on the Gaussian family to model range error and use this to fit parametric surfaces to range data in a maximum-likelihood formulation. Cross-range error is assumed to be negligible and is accommodated by increasing the range error variance.

Each of these recognition papers focus on drawing inferences on the basis of a single range image that represents

distance to the scene as a function of direction from the imaging platform. Further, each addresses measurement errors solely in terms of range uncertainty. In this paper, we address object recognition from 3-D point data, collected from a diverse set of sensor positions, when measurement errors are possible in both range and cross-range.

A number of situations can yield cross-range measurement error, even for a well-calibrated sensing platform. For noncontact measurement systems these can include relative motion between the platform and scene during image acquisition, misregistration when employing multiple range images, and nonzero optical footprint diameter (for example, Verly and Delany¹² report on an experimental ladar system that has a pixel width of 10 cm). A large optical footprint can result in a multipath effect causing a point on the surface to be shifted in the cross-range direction.¹³ For sources of error in contact scanners see Ref. 14. Even neglecting cross-range measurement errors, recognition systems are sensitive to relative errors between the scene and a hypothesized object model, which can arise through uncertainty in the object pose relative to the sensor.

Abstractly we can think of these error sources as resulting in uncertainty in the line-of-sight involved in a particular sample point. We model this uncertainty through a probabilistic relationship between a measured sample point and a corresponding “point of origin” on the object surface. That is, for any point on the object surface there is a probability density over 3-D space governing the measured location of that point. This distribution depends on the object geometry, object pose, and sensor properties, and it will vary from point to point on the object surface. From this model we derive two recognition algorithms. The first is a minimum probability of error classifier that determines the likelihood of a measured point by integrating over all possible origination points on the object’s surface. A somewhat simpler alternative is the generalized likelihood algorithm that eliminates the integration step and instead uses the maximum conditional density from all possible points of origin.

These algorithms reduce to the range error-only model adopted in the previously mentioned papers when the probability mass is restricted to the pixel’s line of sight. A worst-case scenario occurs when the line of sight for a measurement is completely unknown. This situation can arise when processing point data assembled from range images representing multiple perspectives,¹⁵ for example, from mobile platforms or a collection of networked sensors. It is often convenient to represent the data collection as a “point cloud,” an unordered collection of 3-D points measured from a scene without labeling individual points with the direction from which they were measured. We focus on recognition in this extreme case and assume an additive measurement error that follows a spherically symmetric Gaussian distribution with independent observations. We address the recognition performance in terms of sensor network capabilities, including the number of measured points and the measurement noise variance.

In Sec. 2 we discuss the data model adopted to characterize point cloud observations and include the special case of complete uncertainty in line of sight. Recognition algorithms are derived from this model in Sec. 3. Note that, while 3-D point measurement often results in both 3-D

points and 2-D intensity information, we focus on object recognition via the 3-D point data alone. In Sec. 4 we present an analytical evaluation of the classification accuracy when choosing between two spheres of different radii centered at the origin. This simple problem demonstrates that the generalized likelihood test is subject to the peaking phenomenon¹⁶ when measurement noise grows large. In this case, increasing the number of measured points will cause the algorithm performance to deteriorate. Correct classification rates for the minimum probability of error test are shown to monotonically increase with observation size. These results are extended to more general recognition problems in Sec. 5, which contains Monte Carlo simulation results for a three-class problem involving simple geometric shapes. Conclusions follow in Sec. 6.

2 Data Model

Let the set of points comprising the surface of an object be denoted by S . It will be convenient to think of this surface being selected at random from a collection of surfaces \mathcal{S} . Let $\mathbf{X}^* \in S$ be a point on that surface given in 3-D coordinates as $\mathbf{X}^* = [X_1^*, X_2^*, X_3^*]^T$, where a^T denotes the transpose of a . This point is selected from the surface for measurement according to a conditional probability density $p_{\mathbf{X}^*|\theta, \mathcal{S}}(\mathbf{x}^* | \theta, S)$, where θ represents the pose (location and rotation) of the object relative to the measurement platform (and, consequently, the line of sight). In the general case this density varies with θ because geometrical factors such as foreshortening, obscuration, and finite optical footprint make some points more likely to be selected than others. Also, the value of θ will generally vary from point to point within a data collection. This occurs because each pixel in a range image represents a slightly different line of sight to the scene. Moreover, the data may be pooled from multiple range images collected from completely different vantage points.

Let the measured location of the surface point \mathbf{X}^* be denoted by $\mathbf{X} = [X_1, X_2, X_3]^T$. We will refer to \mathbf{X}^* as the point of origin for the measurement \mathbf{X} , and model the observation as $\mathbf{X} = \mathbf{X}^* + \mathbf{N}$, where \mathbf{N} is a multivariate Gaussian distribution with zero mean. That is, $\mathbf{X} \sim \mathcal{N}(\mathbf{X}^*, \Sigma_{\theta, S})$, where $\Sigma_{\theta, S}$ is a 3×3 covariance matrix. In general, the noise variance may depend on the object pose, for instance, there may be a greater variance in the range direction than in the cross-range directions. Noise variance may also depend on the object surface being measured, for instance, very rough surfaces may be more susceptible to multipath effects than smooth objects. The conditional probability density function for \mathbf{X} is

$$p_{\mathbf{X}|\theta, \mathcal{S}, \mathbf{X}^*}(\mathbf{x} | \theta, S, \mathbf{x}^*) = \frac{1}{(2\pi)^{3/2} |\Sigma_{\theta, S}|^{1/2}} \times \exp \left\{ -\frac{1}{2} (\mathbf{x} - \mathbf{x}^*)^T \Sigma_{\theta, S}^{-1} (\mathbf{x} - \mathbf{x}^*) \right\}. \tag{1}$$

Given these conditional densities, the probability density function for an observation \mathbf{X} given that surface S is measured with relative pose θ is

$$p_{\mathbf{X}|\Theta,S}(\mathbf{x}|\theta,S) = \int_{\mathbf{x}^* \in S} p_{\mathbf{X}|\Theta,S,\mathbf{X}^*}(\mathbf{x}|\theta,S,\mathbf{x}^*) \times p_{\mathbf{X}^*|\Theta,S}(\mathbf{x}^*|\theta,S) d\mathbf{x}^*. \quad (2)$$

In keeping with the worst-case scenario discussed in Sec. 1, our examples assume no prior information about the sensor line of sight, and the measured points are modeled as uniformly distributed over the object's surface,

$$p_{\mathbf{X}^*|\Theta,S}(\mathbf{x}^*|\theta,S) = p_{\mathbf{X}^*|S}(\mathbf{x}^*|S) = \frac{1}{\text{Area}(S)}. \quad (3)$$

For this same reason, our examples assume no directional preference in the measurement noise, and Σ is taken to be a diagonal matrix with equal diagonal elements. Finally, we make the approximation that measurement noise does not depend on the surface being interrogated, so the conditional density relating surface points to their measured location $p_{\mathbf{X}|\Theta,S,\mathbf{X}^*}(\mathbf{x}|\theta,S,\mathbf{x}^*) = p_{\mathbf{X}|\mathbf{X}^*}(\mathbf{x}|\mathbf{x}^*)$ is the density function for a spherically symmetric Gaussian distribution with variance σ^2 in all components.

3 Recognition Algorithms

The recognition problem we address is that of determining which object out of a finite set $m \in \{1, 2, \dots, M\}$ resulted in the set of measured surface point locations $\mathcal{X} = \{\mathbf{X}_k\}_{k=1}^K$, which are assumed to be conditionally independent. We assume that the surface S_m of each object is known relative to a fixed world coordinate system. That is, the presence of an object implies a particular shape at a fixed location and orientation. This is a fairly restrictive assumption, and in practice the recognition algorithms of this section must be combined with a search algorithm over the space of possible locations and orientations. Alternatively, if prior distributions for these parameters can be established, then a Bayesian formulation could be adopted to address the uncertainty. Also, we do not make any specific distinction between points on an object and points on the background or other clutter. A fundamental approach would be to use the algorithms to jointly infer the presence of an object and the environment in which it resides (background structure, obscuration, etc.). An alternative approach is to employ a segmentation algorithm, eliminating points that are not likely from one of the objects under consideration.

Given the data model of the previous section, the minimum probability of error decision rule $\phi: \mathbb{R}^{3K} \rightarrow \{1, 2, \dots, M\}$ chooses the object that maximizes the likelihood of the observations. That is,

$$\phi(\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_K) = \underset{m \in \{1, 2, \dots, M\}}{\operatorname{argmax}} P_m \prod_{k=1}^K p_{\mathbf{X}|\Theta,S}(\mathbf{x}_k|\theta_k, S_m), \quad (4)$$

where P_m is the prior probability of observing object m , and $p_{\mathbf{X}|\Theta,S}(\mathbf{x}_k|\theta_k, S_m)$ is from Eq. (2).

The decision rule in Eq. (4) is theoretically optimal, but it has some practical shortcomings. These stem from the need to know the prior distribution of \mathbf{X}^* and to integrate K joint probability densities over the surface of each object under consideration. An alternative decision rule, which is suboptimal but perhaps easier to implement, is the general-

ized likelihood test. The generalized likelihood $\tilde{p}_{\mathbf{X}|\Theta,S}(\mathbf{x}|\theta,S)$ for a surface S given a measured point \mathbf{x} is the maximum conditional probability density of the observation over all points on the surface,

$$\tilde{p}_{\mathbf{X}|\Theta,S}(\mathbf{x}|\theta,S) = \underset{\mathbf{x}^* \in S}{\operatorname{argmax}} p_{\mathbf{X}|\Theta,S,\mathbf{X}^*}(\mathbf{x}|\theta,S,\mathbf{x}^*). \quad (5)$$

The decision rule for the generalized likelihood test is

$$\tilde{\phi}(\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_K) = \underset{m \in \{1, 2, \dots, M\}}{\operatorname{argmax}} \prod_{k=1}^K \tilde{p}_{\mathbf{X}|\Theta,S}(\mathbf{x}_k|\theta_k, S_m). \quad (6)$$

4 Analytical Results

In this section we apply the minimum probability of error and generalized likelihood decision rules to a simple two-class problem that lends itself to straightforward analysis. This will permit us to explore the behavior of the two decision rules with respect to noise level and number of sample points. The problem we address is that of recognizing which of two spheres is present based on some number K of noisy sample points measured independently from the surface of the object in question. The spheres, S_1 and S_2 , have radii $r_1=1$ and $r_2=2$, respectively, and whichever is present is assumed to be centered at the origin of a suitable coordinate system. Measurement noise at each point is trivariate Gaussian with zero mean and diagonal covariance matrix $\Sigma = \sigma^2 \mathbf{I}$, where \mathbf{I} is the 3×3 identity matrix.

Under this noise model, a measurement from a sphere S_r of radius r will result in an observation \mathbf{X} with probability density given by Eq. (2). We first find the density at the point $\mathbf{x} = [0, 0, x_3]^T (x_3 \geq 0)$, for which the integral expression can be simplified and upon conversion from rectangular to polar coordinates becomes

$$\begin{aligned} p_{\mathbf{X}|S}([0, 0, x_3]^T | S_r) &= \frac{1}{8\sqrt{2}\pi^{5/2}\sigma^3} \\ &\times \int_0^{2\pi} \int_0^\pi \exp\left\{-\frac{1}{2\sigma^2}[r^2 \sin^2 \phi + (x_3 - r \cos \phi)^2]\right\} \sin \phi d\phi d\theta \\ &= \frac{1}{4\sqrt{2}\pi^{3/2}rx_3\sigma} [e^{-(x_3-r)^2/2\sigma^2} - e^{-(x_3+r)^2/2\sigma^2}]. \end{aligned}$$

Because of the symmetry of the sphere, the density function must have the same value at all points the same distance from the origin. Thus, the probability density at arbitrary \mathbf{x} can be written as

$$p_{\mathbf{X}|S}(\mathbf{x}|S_r) = \frac{1}{4\sqrt{2}\pi^{3/2}r\|\mathbf{x}\|\sigma} (e^{-\|\mathbf{x}\| - r)^2/2\sigma^2} - e^{-\|\mathbf{x}\| + r)^2/2\sigma^2}). \quad (7)$$

We can use this result to express the probability density function for the radial distance of an observation $P = \|\mathbf{X}\|$ as

$$p_{P|S}(\rho|S_r) = \frac{\rho}{\sqrt{2\pi r\sigma}} (e^{-(\rho-r)^2/2\sigma^2} - e^{-(\rho+r)^2/2\sigma^2}), \quad \text{for } \rho \geq 0. \quad (8)$$

In the subsections below, we find an expression for the mean and variance of the decision rules given that one of the two spheres is being observed. We then fit a Gaussian distribution to the decision rules, an approximation that becomes exact as $K \rightarrow \infty$ by the central limit theorem. The conditional probabilities of error can then be expressed in terms of tail probabilities of a standard Gaussian random variable. We use these expressions to analyze the behavior of the two algorithms as a function of the noise level σ and the sample size K .

4.1 Minimum Probability of Error Analysis

In a two-class problem, the minimum probability of error decision rule in Eq. (4) can be written as a single inequality exploiting the density for ρ derived above. The test becomes

$$\begin{aligned} &\mathcal{H}_1 \\ L(\rho_1, \rho_2, \dots, \rho_K) &\geq 0, \\ &\mathcal{H}_2 \end{aligned} \quad (9)$$

where \mathcal{H}_m is an assertion that the points with distances $\{\rho_k\}_{k=1}^K$ were measured from sphere S_m , and

$$L(\rho_1, \rho_2, \dots, \rho_K) = \sum_{k=1}^K L_k = \sum_{k=1}^K [\ln p_{P|S}(\rho_k|S_1) - \ln p_{P|S}(\rho_k|S_2)]. \quad (10)$$

Because the L_k are independent and identically distributed, $E[L|S_m] = KE[L_k|S_m]$, and $\text{var}(L|S_m) = K\text{var}(L_k|S_m)$. The conditional mean and variance of L_k given S_m can be written as one-dimensional integrals

$$E[L_k|S_m] = \int_0^\infty L_k p_{P|S_m}(\rho|S_m) d\rho, \quad (11)$$

$$\text{var}(L_k|S_m) = \int_0^\infty L_k^2 p_{P|S_m}(\rho|S_m) d\rho - E^2[L_k|S_m]. \quad (12)$$

It is straightforward to evaluate these integrals numerically, since the integrands go to zero quickly as ρ gets large. Figure 1 shows plots of these quantities as a function of σ . For large noise levels, the likelihoods of the two spheres become nearly equal in distribution.

When the object being observed is S_1 , we make a correct decision if $L > 0$. We approximate this conditional probability as

$$\text{Pr}[\text{correct}|S_1] \approx 1 - \Phi\left(-\frac{E[L|S_1]}{\sqrt{\text{var}(L|S_1)}}\right), \quad (13)$$

where Φ is the cumulative distribution function for a standard Gaussian random variable. Similarly, the conditional probability of correctly classifying S_2 is approximated as




Fig. 1 Conditional mean and variance of the minimum probability of error test statistic for the two spheres.

$$\text{Pr}[\text{correct}|S_2] \approx \Phi\left(-\frac{E[L|S_2]}{\sqrt{\text{var}(L|S_2)}}\right). \quad (14)$$

The log-likelihood ratio L is a sum of independent, identically distributed random variables, so for large K these approximations should be quite good. Even for small K the approximation is not unreasonable. For example, with $K = 1$ and $\sigma = 1$ the actual value of $\text{Pr}[\text{correct}|S_1]$ is 0.684, and the approximation above yields 0.719, a difference of 0.035. From Fig. 1, we see that $E[L_k|S_1] > 0$ and $E[L_k|S_2] < 0$, so the conditional probability of making a correct decision is always greater than 1/2. Moreover, increasing the number of sample points will increase the conditional probability of correct recognition, because the conditional mean of L grows linearly with K while the standard deviation grows linearly with \sqrt{K} . For a given sample size K , very large noise levels will cause the conditional probability of correct recognition to approach 1/2 for both spheres.

Figure 2 shows a plot, created using the Gaussian ap-




Fig. 2 Conditional probability of correctly classifying the small sphere S_1 under the minimum probability of error decision rule.




Fig. 3 Difference between the conditional probabilities of correctly classifying the two spheres under the minimum probability of error decision rule.

proximation, of the probability of correctly classifying the small sphere S_1 as a function of noise level for several sample sizes K . The plot shows that for low noise levels, the small sphere can be correctly recognized with near perfect accuracy using only a small number of measured points. The plot also shows graceful degradation in performance as the noise level increases, and it suggests that recognition can still be accurate even within the presence of very large noise levels if more sample points are collected. For instance, the solid curve indicates that if 1000 points are measured, the small sphere is correctly recognized over 97% of the time, even when the noise standard deviation is 3 times as large as the sphere itself. The complementary probability $\Pr[\text{correct}|S_2]$ is similar in behavior but not identical. Figure 3 shows the difference $\Pr[\text{correct}|S_1] - \Pr[\text{correct}|S_2]$ for the same range of noise and sample sizes. The curves show that the minimum probability of error classifier exhibits a slight bias toward the smaller sphere.

4.2 Generalized Likelihood Analysis

For this two class problem, the generalized likelihood decision rule in Eq. (6) can be written as

$$\sum_{k=1}^K [\max_{x_1^* \in S_1} \ln p_{\mathbf{X}|S, \mathbf{X}^*}(\mathbf{x}_k | S_1, \mathbf{x}_1^*) - \max_{x_2^* \in S_2} \ln p_{\mathbf{X}|S, \mathbf{X}^*}(\mathbf{x}_k | S_2, \mathbf{x}_2^*)] \geq 0. \tag{15}$$

\mathcal{H}_1 \mathcal{H}_2

Consider the point $\mathbf{x} = [\rho, 0, 0]^T, \rho > 0$, in which case the maximizers above are $\mathbf{x}_1^* = [r_1, 0, 0]^T$ and $\mathbf{x}_2^* = [r_2, 0, 0]^T$. The maxima of the densities can then be written as

$$\max_{x^* \in S_m} p_{\mathbf{X}|S, \mathbf{X}^*}([\rho, 0, 0]^T, S_m, \mathbf{x}^*) = \frac{1}{(2\pi)^{3/2} \sigma^3} e^{-(\rho - r_m)^2 / 2\sigma^2}.$$

Because of symmetry, the maximum must be the same value for all points at distance ρ from the origin. We can then rewrite Eq. (15) in terms of the distances ρ_k from the observed points to the origin and simplify to obtain

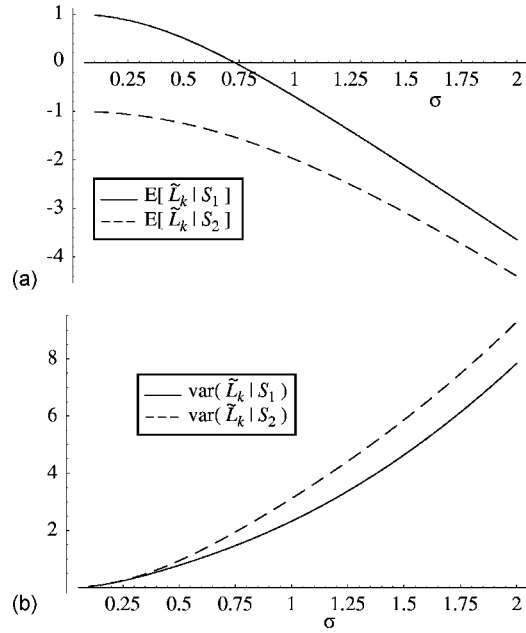


Fig. 4 Conditional mean and variance of the generalized likelihood test statistic for the two spheres.

$$\tilde{L} = \sum_{k=1}^K \tilde{L}_k = \sum_{k=1}^K [(r_2 - r_1)(r_2 + r_1 - 2\rho_k)] \geq 0. \tag{16}$$

\mathcal{H}_1 \mathcal{H}_2

The conditional mean and variance of the terms \tilde{L}_k given S_1 can be found by integrating with respect to ρ using the probability density [Eq. (8)] and are

$$E[\tilde{L}_k | S_1] = (r_2^2 - r_1^2) + 2(r_1 - r_2) \times \left[\frac{\sqrt{2}\sigma}{\sqrt{\pi}} e^{-r_1^2/2\sigma^2} + \frac{r_1^2 + \sigma^2}{r_1} \operatorname{erf}\left(\frac{r_1}{\sqrt{2}\sigma}\right) \right] \tag{17}$$

and

$$\operatorname{var}(\tilde{L}_k | S_1) = 4(r_1 - r_2)^2 \left\{ r_1^2 + \sigma^2 \left(3 - \frac{2}{\pi} e^{-r_1^2/\sigma^2} \right) + \frac{(r_1^2 + \sigma^2) \operatorname{erf}\left(\frac{r_1}{\sqrt{2}\sigma}\right)}{r_1^2} \times \left[-\frac{2\sqrt{2}r_1\sigma}{\sqrt{\pi}} e^{-r_1^2/2\sigma^2} - (r_1^2 + \sigma^2) \operatorname{erf}\left(\frac{r_1}{\sqrt{2}\sigma}\right) \right] \right\}, \tag{18}$$

where $\operatorname{erf}(z) = 2/\sqrt{\pi} \int_0^z e^{-t^2} dt$. Similar expressions hold for the conditional mean and variance given S_2 . Figure 4 shows plots of these quantities versus noise level. Note that the conditional mean given that S_1 is present becomes negative for noise levels above $\sigma \approx 0.7307$. In the limit as $\sigma \rightarrow \infty$, the




Fig. 5 Conditional probability of correctly classifying the small sphere S_1 using the generalized likelihood test.

difference between the conditional means goes to zero. Similarly, the difference between the conditional variances does not get infinitely large but goes to $12-32/\pi \approx 1.814$.

When S_1 is present, we make a correct classification whenever $\tilde{L} > 0$, and when S_2 is present we choose correctly whenever $\tilde{L} < 0$. As before, we approximate these conditional probabilities by fitting a Gaussian distribution,

$$\Pr[\text{correct}|S_1] \approx 1 - \Phi\left(-\frac{E[\tilde{L}|S_1]}{\sqrt{\text{var}(\tilde{L}|S_1)}}\right), \quad (19)$$

$$\Pr[\text{correct}|S_2] \approx \Phi\left(-\frac{E[\tilde{L}|S_2]}{\sqrt{\text{var}(\tilde{L}|S_2)}}\right). \quad (20)$$

Again, we expect this to be a reasonable approximation, even for small K . For example, Monte Carlo simulation with $K=1$ and $\sigma=1$ yields an estimate for $\Pr[\text{correct}|S_1]$ of 0.351, and the approximation above yields 0.382, a difference of 0.031. From Fig. 4 we see that for small noise levels, $E[\tilde{L}|S_1] > 0$ and $E[\tilde{L}|S_2] < 0$. Consequently we expect that the conditional probabilities of correctly classifying each sphere will be greater than 1/2, and increasing the number K of sample points will increase these probabilities. However, for large noise levels the situation is quite different. Beyond the critical value $\sigma \approx 0.7307$, $E[\tilde{L}|S_1] < 0$, and the probability of correctly recognizing S_1 when it is present will be less than 1/2. Moreover, increasing the number of sample points will *decrease* the conditional mean of \tilde{L} and will result in *reduced* classification accuracy for S_1 .

This is demonstrated in Fig. 5, which shows the conditional probability of correctly classifying S_1 with the generalized likelihood test as a function of the noise level for several sample sizes K . The plot shows that for small noise levels, nearly perfect classification of the small sphere S_1 can be achieved with only a small number of points. However, catastrophic failure ensues when the noise level surpasses the critical value $\sigma \approx 0.7307$. Beyond this value, increasing the number of sample points actually reduces the probability of correctly classifying S_1 . By contrast, the probability of correctly classifying the larger sphere S_2 is




Fig. 6 Overall probability of correct classification as a function of sample size K using the generalized likelihood test.

nearly one for all values of σ when more than one or two sample points are available. As a result, the overall probability of correct classification, $\Pr[\text{correct}] = 0.5(\Pr[\text{correct}|S_1] + \Pr[\text{correct}|S_2])$, exhibits a peak for some value of $K < \infty$ and decreases beyond this value. Figure 6 shows the overall probability of correct recognition as a function of K for values of σ slightly below and slightly above the critical value. For $\sigma=0.7$, overall accuracy is seen to increase monotonically with K , while for $\sigma=0.8$ a clear peak is observed around $K=4$ with decreasing performance for larger K . This behavior, commonly referred to as the “peaking phenomenon,” is most frequently reported with respect to adaptive classifiers designed on the basis of training data and incorporating some form of feature selection.¹⁷ In the shape recognition example of this section, we observe the phenomenon for a slightly different class of suboptimal decision rule, one that involves neither model training nor feature ordering and selection.

5 Experimental Results

The previous section contains an analytical treatment of the performance delivered by the two decision rules when the objects involved have the same shape but different scales. This section contains Monte Carlo simulation results demonstrating the performance behavior when the objects have the same scale but different shape, and these results suggest that many of the behavioral phenomena derived for the two sphere example are also evident in more general problems. The tests are applied to a problem involving three simple geometrical shapes, which are equally likely *a priori*. The objects each encompass 1 cubic unit of volume and are:

- a sphere centered at the origin with radius $(3/4\pi)^{1/3}$;
- a cube centered at the origin with edge length 1 and vertices at $[\pm 1/2, \pm 1/2, \pm 1/2]$; and
- a tetrahedron with edge length $a=(12/\sqrt{2})^{1/3}$ and vertices at $[0, 0, \sqrt{6}/4a]$, $[0, \sqrt{3}/3a, -\sqrt{6}/12a]$, $[-a/2, -\sqrt{3}/6a, -\sqrt{6}/12a]$, and $[a/2, -\sqrt{3}/6a, -\sqrt{6}/12a]$.

For purposes of simulation, points are selected for measurement from a uniform distribution over the surface of each object. As before, we model the measurement noise as independent, spherically symmetric Gaussian, $N \sim \mathcal{N}(0, \Sigma)$, where the covariance matrix $\Sigma = \sigma^2 \mathbf{I}$, and \mathbf{I} is the 3×3 identity matrix. We conduct two sequences of experiments, one with $K=10$ measured points and the other with K




Fig. 7 Probability of correct classification for both decision rules with $K=10$ and $K=100$ measured points, as a function of noise level.

$=100$, evaluating algorithm performance at noise levels $\sigma = 0.1, 0.2, \dots, 0.9$. Performance is assessed for each object as the fraction of 1000 trials for which the object is classified correctly.

Figure 7 shows the overall performance for both decision rules and for both sample sizes, as a function of noise level. The left two bars in each group show the performance of the generalized likelihood test and the minimum probability of error (optimal) test with $K=10$ sample points, respectively. The right two bars in each group show the performance of these tests with $K=100$ sample points. As expected, the minimum probability of error test gives a higher overall probability of correct classification than the generalized likelihood test. For small noise levels and large sample sizes the generalized likelihood test delivers a very high correct classification rate. However, as the noise level increases, the incremental benefit of large sample sizes begins to decrease, and beyond $\sigma \approx 0.7$ the generalized likelihood rule with $K=100$ points actually performs worse than with $K=10$.

Figure 8 shows a performance breakdown by object for both tests and both sample sizes. In each panel, the black bar represents $\Pr[\text{correct}|\text{sphere}]$, the gray bar represents $\Pr[\text{correct}|\text{cube}]$, and the white bar represents $\Pr[\text{correct}|\text{tetrahedron}]$. Figure 8(a), which contains results for the generalized likelihood test with $K=10$ measured points, shows that the probability of correctly recognizing the sphere decreases most rapidly with increasing noise level, followed by the probability of correctly recognizing the cube. The probability of correctly recognizing the tetrahedron decreases initially, but levels out around $\sigma=0.5$.

In comparing these results to Fig. 8(b), which shows performance for $K=100$, we see details of the catastrophic behavior exhibited in the previous section. For $\sigma=0.1$, increasing sample size results in increasing conditional probabilities of correct recognition for all three objects, and the accuracy is very good. For $\sigma=0.3$, increasing from $K=10$ to $K=100$ sample points increases the probability of correctly recognizing the cube and the tetrahedron, however it results in a zero probability of correctly recognizing the sphere. By $\sigma=0.5$, the probability of correctly recognizing the cube begins to diminish with increasing sample sizes, and beyond $\sigma=0.9$ the decision rule selects the tetrahedron in every trial, regardless of which object actually generated the sample. Note that while each object encloses the same area, the tetrahedron has the largest surface area, and the cube has the second largest. Thus, these results are consistent with the previous section and suggest that the generalized likelihood test is biased toward objects with large surface area.

Figures 8(c) and 8(d) show performance under the minimum probability of error (optimal) decision rule with $K=10$ and $K=100$, respectively. The graphs show a gradual decline in the conditional correct recognition rates for all objects as the noise level increases. For each object and for each noise level, increasing the number of sample points increases probability of correct recognition. The correct recognition rate for the tetrahedron is consistently greater than for the other objects. The slight bias of the minimum probability of error rule toward the object with smaller surface area that was observed in the two sphere problem is not a generally observable property, as it is evident that the

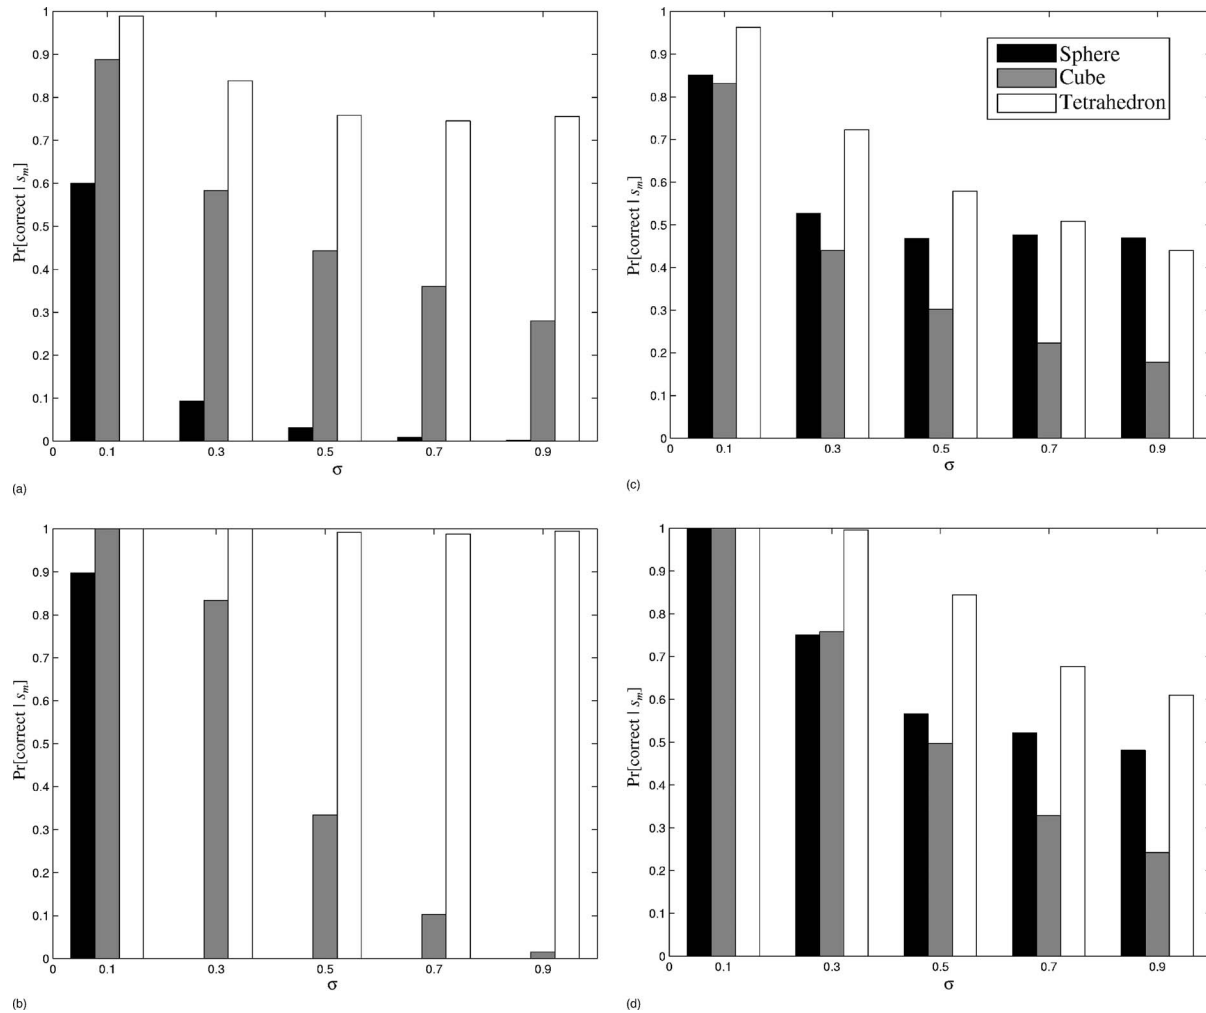


Fig. 8 Probability of correct classification for both decision rules with $K=10$ and $K=100$ measured points, as a function of noise level. (a) Generalized likelihood, $K=10$, (b) generalized likelihood, $K=100$, (c) optimal decision rule, $K=10$, and (d) optimal decision rule, $K=100$.

correct recognition rate for the cube is generally lower than for the sphere. The results of this experiment are generally consistent with the directed Hausdorff distances between the objects. The directed Hausdorff distance from a set of 3-D points A to another set B is the quantity $\sup_{a \in A} \inf_{b \in B} d(a, b)$, where $d(a, b)$ denotes the Euclidian distance between points a and b . The set of points comprising the tetrahedron has a larger directed distance to the other objects, and the tetrahedron has a higher probability of correct classification. The directed distance from the cube to the other two objects is relatively small, and the cube is frequently classified as either a sphere or a tetrahedron. While the directed distance from the sphere to the cube is not large, the distance from the sphere to the tetrahedron is, and the sphere has a correct recognition rate that falls between the other two.

6 Conclusions

We have formulated a general model for noisy measurements of 3-D points from the surface of an object, which accounts for both range and cross-range measurement error. This model was used to construct two decision rules for

choosing the object that yielded a given collection of measured points. The model and algorithms build upon earlier work by a variety of researchers, which assumes that range error dominates the measurements and that cross-range error is negligible. The (optimal) minimum probability of error decision rule forms the likelihood of each object, accommodating uncertainty in the surface point that yielded a given measured location by integrating the joint density of the pair over the entire surface of the object. In contrast, the (suboptimal) generalized likelihood decision rule finds the single best point on the object surface for each measured location. The behavior of these two decision rules was compared in an analytically tractable problem involving two objects with the same shape but different sizes. Further analysis was performed by encoding the decision rules into a pair of recognition algorithms, using Monte Carlo simulation to estimate their performance in a problem with three objects of the same size but different shapes.

For very low noise levels, the performance of the generalized likelihood test is comparable to that of the minimum probability of error test, though there is a slight bias toward objects with large surface area. However, there

tends to be a critical value of the noise level, beyond which the generalized likelihood decision rule exhibits a catastrophic failure mode. This critical noise level varies with the size and shape of each object relative to the other objects under consideration. If the noise level exceeds this critical value for a given object, the probability of correctly recognizing the object drops dramatically, and increasing the number of measured points on its surface actually decreases the recognition accuracy. In contrast, the minimum probability of error decision rule can deliver good performance even at high noise levels, provided that large sample sizes are available.

The performance properties of the minimum probability of error rule comes at some implementation burden. For complicated shapes, implementation of the minimum probability of error rule entails significant computational complexity. This is due to the required integration over the set of points on the object's surface. For polygonal CAD models the integration can be approximated by a summation over each polygon, an operation which is straightforward but potentially time consuming for models with a high level of detail. On the other hand, the generalized likelihood test has a relatively simple implementation, requiring a maximization over the object's surface rather than an integration. For spherically symmetric noise distribution, this maximum will occur at the closest point on the object to the measured value. If an object is represented by a triangulated CAD model of sufficient resolution, this maximum may be approximated as the CAD model vertex closest to the measured value. Simple hashing techniques can potentially be used to greatly reduce the cost of searching for this vertex.

Acknowledgments


This work was supported by the Office of Naval Research grant N00014-03-1-0110.

References

1. H. Haitjema, W. Pril, and P. H. J. Schellekens, "A silicon-etched probe for 3-D coordinate measurements with and uncertainty below $0.1 \mu\text{m}$," *IEEE Trans. Instrum. Meas.* **50**(6), 1519–1523 (Dec. 2001).
2. J. Li, Z.-S. Liu, and P. Stoica, "3-D target feature extraction via interferometric SAR," *IEE Proc. Radar, Sonar, and Navigation* **144**(2), 71–80 (Apr. 1997).
3. M. Sainz, R. Pajarola, A. Mercade, and A. Susin, "A simple approach for point-based object capturing and rendering," *IEEE Comput. Graphics Appl.* **24**(4), 24–33 (Jul. 2004).
4. J. J. Le Moigne and A. M. Waxman, "Structured light patterns for robot mobility," *IEEE J. Rob. Autom.* **4**(5), 541–548 (Oct. 1988).
5. C. Reich, R. Ritter, and J. Thesing, "3-D shape measurement of complex objects by combining photogrammetry and fringe projection," *Opt. Eng.* **39**(1), 224–231 (Jan. 2000).
6. P. J. Besl and R. C. Jain, "Three-dimensional object recognition," *ACM Comput. Surv.* **17**(1), 75–145 (Mar. 1985).
7. Q. Zheng, S. Z. Der, and H. I. Mahmoud, "Model-based target recognition in pulsed lidar imagery," *IEEE Trans. Image Process.* **10**(4), 565–572 (Apr. 2001).
8. B. A. Hutchinson, R. L. Galbraith, B. L. Stann, and S. Z. Der, "Simulation-based analysis of range and cross-range resolution requirements for the identification of vehicles in lidar imagery," *Opt. Eng.* **42**(9), 2734–2745 (Sep. 2003).
9. T. J. Green and J. H. Shapiro, "Detecting objects in three-dimensional laser radar range images," *Opt. Eng.* **33**(4), 865–874 (Mar. 1994).
10. A. D. Lanterman, "Jump-diffusion algorithm for multiple target recognition using laser radar range data," *Opt. Eng.* **40**(8), 1724–1728 (Aug. 2001).
11. R. T. Whitaker and J. Gregor, "A maximum-likelihood surface estimator for dense range data," *IEEE Trans. Pattern Anal. Mach. Intell.* **24**(10), 1372–1387 (Oct. 2002).
12. J. G. Verly and R. L. Delanoy, "Model-based automatic target recognition (ATR) system for forwardlooking groundbased and airborne imaging laser radars (LADAR)," *Proc. IEEE* **84**(2), 126–163 (Feb. 1996).
13. M. D. Adams, "Lidar design, use, and calibration concepts for correct environmental detection," *IEEE Trans. Rob. Autom.* **16**(6), 753–761 (Dec. 2000).
14. R. P. Johnson, Q. Yang, and C. Butler, "Dynamic error characteristics of touch trigger probes fitted to coordinate measurement machines," *IEEE Trans. Instrum. Meas.* **47**(5), 1168–1172 (Oct. 1998).
15. G. Blais and M. D. Levine, "Registering multiview range data to create 3D computer objects," *IEEE Trans. Pattern Anal. Mach. Intell.* **17**(8), 820–824 (Aug. 1995).
16. S. J. Raudys and P. Vitalijus, "On dimensionality, sample size, classification error, and complexity of classification algorithm in pattern recognition," *IEEE Trans. Pattern Anal. Mach. Intell.* **PAMI-2**(3), 242–252 (May 1980).
17. A. Jain and D. Zongker, "Feature selection: evaluation, application, and small sample performance," *IEEE Trans. Pattern Anal. Mach. Intell.* **19**(2), 153–158 (Feb. 1997).



Xin Zhou received his BS degree from the University of Science and Technology of China and his MS degree from Chinese Academy of Sciences, in 2000 and 2003, respectively. He is currently studying in the Department of Systems and Information Engineering at the University of Virginia for his PhD degree. His research interests include pattern recognition, computer vision, and robotics.



Michael D. DeVore received BS degrees in electrical engineering, computer engineering, and mathematics from the University of Missouri at Columbia in 1991 and an MS degree in electrical engineering in 1993. From 1993 through 1998 he was responsible for system integrity and integration and for operations support at Amdocs, Inc., a solutions provider to the telecommunications industry. In 2001 he received the DSc degree in electrical engineering at Washington University in St. Louis, with dissertation research in automatic target recognition from synthetic aperture radar data. Dr. DeVore is currently an assistant professor in the Systems and Information Engineering Department at the University of Virginia. His research interests include target recognition, information systems for computer vision, shape-based object recognition, statistical model assessment, model representation and storage, and performance of recognition systems under dynamic resource constraints. Dr. DeVore is a member of IEEE, SPIE, and the Optical Society of America.