Estimating Surface Orientation from Sonar Images

Nicholas C Makris

Naval Research Laboratory
Washington, DC 20375
USA
Email: makris@abyss.nrl.navy.mil

Abstract

A maximum likelihood method for estimating remote surface orientation from multi-static sonar images is presented. It is assumed that the sonar images are corrupted by signal-dependent noise, known as speckle, arising from complex Gaussian field fluctuations, and that the surface reflectance properties are effectively Lambertian. The minimum number of independent samples necessary for maximum likelihood estimates to be asymptotically unbiased and attain classical estimation theory's lower bound on resolution are also derived.

1. Introduction

Sonar images of remote surfaces are typically corrupted by signal-dependent noise known as speckle. This noise arises when wavelength scale roughness on the surface causes a random interference pattern in the sound field scattered from it by an active system. Relative motion between source, surface and receiver causes the received field to fluctuate over time with circular complex Gaussian random (CCGR) statistics [1]. Underlying these fluctuations, however, is the expected radiant intensity from the surface, from which its orientation may be inferred. In many cases of practical importance, Lambert's Law is appropriate for such inference because variations in the projected area of a surface patch, as a function of source and receiver orientation, often cause the predominant variations in its radiance. Maximum likelihood estimators for Lambertian surface orientation are derived. These are asymptotically optimal when a sufficiently large number of independent samples is available, even though the relationship between surface orientation and measured radiance is generally nonlinear. Here the term optimal means that the estimate is unbiased and its mean-square error attains classical estimation theory's lower bound, the inverse Fisher information matrix, which is also derived by analytic means. By evaluating this bound, it is found that optimal resolution varies significantly with illumination direction and measurement diversity. In a particularly compelling example, it is shown that the minimum error in estimating the angle of incidence with respect to a Lambertian surface is at best proportional to the *cotangent* of this angle, so that surface orientation varies from irresolvable at normal incidence to perfectly resolvable at shallowest grazing. General expressions for the requisite number of independent samples necessary for asymptotic optimality of the maximum likelihood estimate are derived. When evaluated, this number is also found to vary significantly with illumination direction and measurement diversity, as may be expected from the inherently nonlinear nature of the surface estimation problem. A far more detailed account of this material will appear in Reference 2.

2. Radiometry

The flux $d\Phi$, received in a sonar beam of solid angle $d\beta$, is related to the area of the resolved surface patch dA_β, the local surface radiance L_β, and the solid angle subtended by the sonar aperture $d\Omega$, by the linear equation

$$d\Phi = dA_\beta L_\beta d\Omega \tag{1}$$

The solid angle subtended by the sonar aperture, from the surface patch dA_β, is $d\Omega = \cos\psi_r \, dA/r^2$, where dA is the area of the aperture, $\cos\psi_r$ is the foreshortening of the surface patch with respect to the receiver, ψ_r is the scattering angle, and r is the range to the aperture. The intensity of the received beam is then

$$I_\beta = \frac{d\Phi}{dA} = dA_\beta L_\beta \frac{\cos\psi_r}{r^2}. \tag{2}$$

Assuming that the sonar is of sufficiently high resolution that it resolves an elemental surface patch dA_β that is locally planar and small enough so that

$$d\beta = dA_\beta \frac{\cos\psi_r}{r^2}, \tag{3}$$

surface radiance can be directly measured by the sonar as

$$\frac{dI_\beta}{d\beta} = L_\beta . \tag{4}$$

For a Lambertian surface,

$$L_\beta = \rho E \cos\psi_i , \tag{5}$$

so that the radiance measured in a sonar image of the scene L_β is independent of the viewing direction ψ_r. It follows a linear relationship with the foreshortening $\cos\psi_i$ of the surface patch, the surface irradiance E, and the surface bidirectional reflectance distribution function ρ which is $1/\pi$ for a perfectly reflecting Lambertian surface. Here ψ_i is the angle of incident insonification and E is defined as the incident flux per unit area on the surface of albedo $\pi\rho$.

3. Measurement Statistics

Let the stochastic measurement vector $\boldsymbol{R}$ contain the independent statistics R_k whose expected values $\sigma_k(\boldsymbol{a})=\langle R_k\rangle$ are linearly related to measured surface radiance for $k=1,2,3 \dots N$, where the vector $\boldsymbol{a}$ contains the surface orientation parameters a_j to be estimated from the measurements $\boldsymbol{R}$ for $j=1,2,3 \dots N_a$. More succinctly, let $\boldsymbol{\sigma}(\boldsymbol{a}) = \langle \boldsymbol{R} \rangle$.

Assuming the R_k are corrupted by CCGR field fluctuations, the conditional probability distribution for the measurements $\boldsymbol{R}$ given parameter vector $\boldsymbol{a}$ is the product of gamma distributions [1][3]

$$P_R(\boldsymbol{R}|\boldsymbol{a}) = \prod_{k=1}^{N} \frac{\left(\frac{\mu_k}{\sigma_k(\boldsymbol{a})}\right)^{\mu_k} (R_k)^{\mu_k - 1} \exp\left\{-\mu_k \frac{R_k}{\sigma_k(\boldsymbol{a})}\right\}}{\Gamma(\mu_k)} . \tag{6}$$

The quantity μ_k is the number of coherence cells in the measurement average used to obtain R_k [1][3]. This number is equal to the signal-to-noise ratio (SNR) $\langle R_k\rangle^2 / \left(\langle R_k^2\rangle - \langle R_k\rangle^2\right)$. For example, μ_k equals the time-bandwidth product of the received field if each R_k is obtained from a continuous but finite-time average. Additionally, μ_k can be interpreted as the number of stationary speckles averaged over a finite spatial aperture in the image plane or the number of stationary multi-look images averaged for a particular scene [3].

4. Classical Estimation Theory

The maximum likelihood estimator $\hat{a}$, for the parameter set a, maximizes the likelihood function for a. The likelihood function for a is the conditional probability $P_R(R|a)$ evaluated at the measured values R. If $\ln P_R(R|a)$ is differentiable, the maximum likelihood estimate $\hat{a}$ can be found by solving the equation

$$\left.\frac{\partial \ln P_R(R|a)}{\partial a}\right|_{a=\hat{a}} = 0. \tag{7}$$

For sufficiently large μ_k, the maximum likelihood estimate asymptotically obeys the Gaussian distribution

$$P_{\hat{a}}(\hat{a}|a) \approx \sqrt{\frac{|J|}{(2\pi)^{N_a}}} \exp\left\{-\frac{1}{2}(\hat{a}-a)^{\mathrm{T}} J(\hat{a}-a)\right\}, \tag{8}$$

where J is the Fisher information matrix, with elements

$$J_{ij}(a) = \sum_{k=1}^{N}\left(\mu_k \frac{\partial \ln \sigma_k(a)}{\partial a_i}\frac{\partial \ln \sigma_k(a)}{\partial a_j}\right), \tag{9}$$

following References [1][3]. In this limit, the estimate $\hat{a}$ is unbiased and attains classical estimations theory's lower bound J^{-1} on the mean square error of any unbiased estimate of a. In the deterministic limit $\mu_k \to \infty$, where the R_k are obtained from exhaustive sample averages, $P_{\hat{a}}(\hat{a}|a)$ becomes the delta function $\delta(\hat{a}-a)$.

5. A Higher Order Asymptotic Approach to Inference

I have recently developed higher order asymptotic methods for determining the sample size μ_k necessary for the maximum likelihood estimate to be effectively unbiased and attain the classical bound on mean-square error. These are derived in detail in Reference 2. Assuming uniform sampling, $\mu=\mu_k$, and letting $g(\sigma) = \hat{a}(\langle R\rangle)$, the bias of $\hat{a}_i$ will be negligible when

$$\mu >> \left|\frac{\sum_{k=1}^{N}\sigma_k^2 \frac{\partial^2 g_i}{\partial \sigma_k^2}}{2g_i}\right|, \tag{10}$$

and the mean-square error covariance $\left\langle(\hat{a}_i - a_i)^2\right\rangle$ will attain the lower bound $\left[J^{-1}\right]_{ii}$ when

$$\mu >> \left|\frac{\sum_{k=1}^{N} 2\sigma_k^3 \frac{\partial g_i}{\partial \sigma_k}\frac{\partial^2 g_i}{\partial \sigma_k^2} + \frac{1}{4}\sum_{1=1}^{N}\sum_{m=1}^{N}(1-\delta_{1m})\sigma_1^2\sigma_m^2\left\{\frac{\partial^2 g_i}{\partial \sigma_1^2}\frac{\partial^2 g_i}{\partial \sigma_m^2} + 2\left(\frac{\partial^2 g_i}{\partial \sigma_1 \partial \sigma_m}\right)^2\right\}}{\sum_{k=1}^{N}\sigma_k^2\left(\frac{\partial g_i}{\partial \sigma_k}\right)^2}\right|, \tag{11}$$

where the denominator of this expression merely equals $\left[J^{-1}\right]_{ii}$.

6. Inferring Lambertian Surface Orientation

6.1 PROBLEM STATEMENT

For measurement k, a collimated source with known unit incident direction s_k irradiates a planar Lambertian surface with unknown unit normal vector $\boldsymbol{n}$. For each measurement, the receiver measures Lambertian surface radiance from any hemispherical observation position within view of the surface. For convenience, a Cartesian coordinate system (x,y,z) is adopted where the ultimate viewing direction is aligned with the positive z axis. Because surface irradiance E_k is presumed known, given knowledge of source power, directionality, and transmission characteristics to the surface, it is deterministically scaled out of the measured surface radiance leaving $\sigma_{L_k}=\langle L_k\rangle/E_k$. When signal-independent additive CCGR noise of intensity $\sigma_{N_k} d\beta_k E_k$ is also measured with the radiant field from the surface, the expected measurement vector $\langle \boldsymbol{R} \rangle$ becomes $\sigma(\boldsymbol{a}) = \sigma_L(\boldsymbol{a}) + \sigma_N(\boldsymbol{a})$.

Lambert's Law for the expected radiometric component of the data is then

$$\sigma_L = \mathbf{S}\boldsymbol{x}, \tag{12}$$

where the matrix $\mathbf{S}$ is defined by

$$\mathbf{S}^{\mathrm{T}} = \left[s_1 \; s_2 \; s_3 \; \ldots s_N \right], \tag{13}$$

and the vector $\boldsymbol{x}$ equals $\rho\boldsymbol{n}$.

The general problem is to determine both the Labertian surface normal vector $\boldsymbol{n}$ and the albedo $\pi\rho$ from the fluctuating measurements $\boldsymbol{R}$. The surface normal is typically expressed in terms of the surface gradient components,

$$\boldsymbol{n}^{\mathrm{T}} = \left[-p_n \;\; -q_n \;\; 1 \right] / \left(1 + p_n^2 + q_n^2 \right)^{1/2}, \tag{14}$$

where

$$p_n = \frac{\partial z}{\partial x}, \; q_n = \frac{\partial z}{\partial y}, \tag{15}$$

or in terms of spherical coordinates

$$\boldsymbol{n}^{\mathrm{T}} = \left[\cos\phi_n \sin\theta_n \;\; \sin\phi_n \sin\theta_n \;\; \cos\theta_n \right]. \tag{16}$$

6.2 THE ANGLE OF INCIDENCE

Suppose that the angle of incidence ψ is to be estimated from a single measurement R, with variance σ^2/μ, given that the albedo $\rho\pi$ is known. From (9), the resulting mean-square error bound is

$$E\left[(\tilde{\psi} - \psi)^2 \right] \geq J^{-1} = \frac{(\cos\psi + \sigma_N)^2}{\mu \sin^2 \psi}, \tag{17}$$

for any unbiased estimate $\tilde{\psi}$, which becomes

$$E\left[(\tilde{\psi} - \psi)^2 \right] \geq J^{-1} = \frac{\cot^2 \psi}{\mu}, \tag{18}$$

when the signal-independent noise is negligible. These expressions show resolution of the incident angle to be highest when the Lambertian surface is illuminated at shallow grazing and lowest when the surface is illuminated near normal incidence. This can be motivated physically by noting that for shallow grazing angle illumination Lambert's Law has a first order dependence that is proportional to the incident angle. Conversely, for illumination near normal incidence Lambert's Law is independent of the incident angle to first order. It is also significant that when the rms-error bound is finite, it can be reduced in proportion to the square-root of the number of independent samples μ averaged to obtain the radiometric statistic R.

The maximum likelihood estimate for the the angle of incidence is

$$\hat{\psi} = \cos^{-1}\left(\frac{R - \sigma_N}{\rho}\right) \tag{19}$$

Many of the potential benefits and difficulties associated with maximum likelihood estimation can be illustrated by examining the statistical properties of $\hat{\psi}$.

For the remainder of this section, let σ_N be negligible, as may be expected in practical imaging systems except at shallow grazing where ψ is very near $\pi/2$. First of all, because R is a gamma variate and can take on any positive definite value, the estimate $\hat{\psi}$ is real for $0 \le R/\rho \le 1$ and imaginary for $R/\rho > 1$. The probability that $\hat{\psi}$ is real is found to be $\gamma(\mu, \mu/\cos\psi)/\Gamma(\mu)$ by appropriately integrating $P_R(R|\psi)$. But this leaves finite probability $\Gamma(\mu, \mu/\cos\psi)/\Gamma(\mu)$ that $\hat{\psi}$ is imaginary. More specifically, the statistic $\hat{\psi}$ is distributed according to

$$P_{\hat{\psi}}(\hat{\psi}|\psi) = \rho \sin\hat{\psi} P_R(\rho\cos\hat{\psi}|\psi) \qquad \text{over } 0 \le \hat{\psi} \le \pi/2, \tag{20}$$

for $\tilde{\psi}$ real, and

$$P_{\hat{\psi}}(\hat{\psi}|\psi) = \rho \sinh\hat{\psi} P_R(\rho\cosh\hat{\psi}|\psi) \qquad \text{over } 0 \le \hat{\psi} \le \infty, \tag{21}$$

for $\hat{\psi}$ imaginary. The probability $\Gamma(\mu, \mu/\cos\psi)/\Gamma(\mu)$ that $\hat{\psi}$ is imaginary decreases as the angle of incidence ψ and the sample size μ increase, as does the bias of $\hat{\psi}$.

Apparently, CCGR fluctuations in the radiant field can lead to unphysical maximum likelihood estimates of the incident angle ψ. This can be remedied by reconditioning the maximum likelihood estimate, given ancillary information [4] that is $\hat{\psi}$ real, so that

$$P_{\hat{\psi}}\left(\hat{\psi}\middle|\psi, \hat{\psi} = \mathrm{Re}\{\hat{\psi}\}\right) = \rho\sin\hat{\psi}P_R\left(\rho\cos\hat{\psi}\middle|\psi\right)\Gamma(\mu)/\gamma(\mu, \mu/\cos\psi) \qquad \text{over } 0 \le \hat{\psi} \le \pi/2. \tag{22}$$

For sufficiently large samples μ, the relationship between maximum likelihood estimate $\hat{\psi}$ and data R approaches linearity, so that $\hat{\psi}$ obeys the Gaussian distribution

$$P(\hat{\psi}|\psi) \approx \sqrt{\frac{\mu}{2\pi\cot^2\psi}}\exp\left(-\frac{1}{2}\mu\frac{\{\hat{\psi}-\psi\}^2}{\cot^2\psi}\right), \tag{23}$$

with bias vanishing and variance equaling the inverse Fisher information.

Following Section 5, $\hat{\psi}$ is effectively unbiased when

$$\mu >> \left|\frac{\cot^3 \psi}{2\psi}\right|, \tag{24}$$

and effectively attains the bound J^{-1} when

$$\mu >> 2\cot^2 \psi . \tag{25}$$

As these expressions show, the number of samples necessary for $\hat{\psi}$ to behave as a minimum variance unbiased estimate varies nonlinearly from unity at shallow grazing angles to an arbitrarily large number near normal incidence.

Even if the location of the surface is known with respect to the source of illumination, knowledge of the angle of incidence only places the surface within a cone about the incident vector s. At minimum, a second measurement is necessary to determine the orientation of the surface in this cone, and a third is needed to determine the albedo.

6.3 SURFACE ORIENTATION AND ALBEDO

The 3-D parameter vector $\boldsymbol{x}$ is to be estimated from the potentially over-determined N-D measurement vector $\boldsymbol{R}$. From (9) and (12), the mean-square error bound on any unbiased estimate $\tilde{\boldsymbol{x}}$ is

$$E\left[\left(\tilde{\boldsymbol{x}}-\boldsymbol{x}\right)\left(\tilde{\boldsymbol{x}}-\boldsymbol{x}\right)^{\mathrm{T}}\right] \geq \boldsymbol{J}_x^{-1} = \left[\mathrm{S}^{\mathrm{T}} \boldsymbol{J}_{\sigma} \mathrm{S}\right]^{-1}, \tag{26}$$

where $\left[\boldsymbol{J}_{\sigma}\right]_{ij} = \delta_{ij}\mu_{ij} / \sigma_i^2$ is infinite when all incident vectors s_k are tangent to $\boldsymbol{n}$ in the absence of signal-independent noise.

More generally, such as for surface gradient $\boldsymbol{a}^{\mathrm{T}}=[p\ q\ r]$, or polar coordinate $\boldsymbol{a}^{\mathrm{T}}=[\theta\ \phi\ \rho]$ parametrizations, the bound

$$E\left[(\tilde{\boldsymbol{a}} - \boldsymbol{a})(\tilde{\boldsymbol{a}} - \boldsymbol{a})^{\mathrm{T}}\right] \geq \boldsymbol{J}_a^{-1} = \frac{\partial \boldsymbol{a}}{\partial \boldsymbol{x}} \boldsymbol{J}_x^{-1} \frac{\partial \boldsymbol{a}}{\partial \boldsymbol{x}}^{\mathrm{T}} \tag{27}$$

becomes singular when all the s_k are coplanar but not tangent to the surface for non-zero σ_N, or when the Jacobian $\left|\frac{\partial \boldsymbol{a}}{\partial \boldsymbol{x}}\right|$ is singular. When σ_N vanishes, $\left|\boldsymbol{J_x}\right|^{1/2}\left|\boldsymbol{J}_{\sigma}\right|^{-1/2}$ can be interpreted as the effective weighted volume of incident vectors s_k. For example, when $\boldsymbol{R}$ is a 3-D vector, the bound is

$$\left[\boldsymbol{J}_x^{-1}\right]_{ij} = \frac{\sigma_1^2\left[s_2 \times s_3\right]_i\left[s_2 \times s_3\right]_j + \sigma_2^2\left[s_3 \times s_1\right]_i\left[s_3 \times s_1\right]_j + \sigma_3^2\left[s_1 \times s_2\right]_i\left[s_1 \times s_2\right]_j}{\left(s_1 \cdot s_2 \times s_3\right)^2}, \tag{28}$$

so that $\left|\boldsymbol{J_x}\right|^{1/2}\left|\boldsymbol{J}_{\sigma}\right|^{-1/2}$ simply is the volume $\left(s_1 \cdot s_2 \times s_3\right)$ of the parellelepiped of incident vectors. Behavior of the Jacobian $\left|\frac{\partial \boldsymbol{a}}{\partial \boldsymbol{x}}\right|$ depends upon the final coordinates $\boldsymbol{a}$, as can be seen from the respective forms $\left(p^2+q^2+1\right)^{3/2}/\rho^2$ and $1/\left(\rho^2 \sin^2\theta\right)$ for the surface gradient and polar systems.

The maximum likelihood estimate, given linearity between $\boldsymbol{x}$ and σ,

$$\hat{\boldsymbol{x}} = \left[\mathrm{S}^{\mathrm{T}} \boldsymbol{J}_{\sigma} \mathrm{S}\right]^{-1} \mathrm{S}^{\mathrm{T}} \boldsymbol{J}_{\sigma}\left(\boldsymbol{R} - \sigma_N\right), \tag{29}$$

is unbiased and attains the error bound $J_{\boldsymbol{x}}^{-1}$.

Given this information, the maximum likelihood estimates for albedo $\pi\hat{\rho} = \pi\left|\hat{\boldsymbol{x}}\right|$, surface normal $\boldsymbol{n} = \hat{\boldsymbol{x}}/\hat{\rho}$, cone and polar angles $\hat{\theta} = \cos^{-1}\hat{n}_3, \hat{\varphi}_n = \tan^{-1}\dfrac{\hat{n}_2}{\hat{n}_1}$, and surface gradient components $\hat{p}_n = -\dfrac{\hat{n}_1}{\hat{n}_3}, \hat{q}_n = -\dfrac{\hat{n}_2}{\hat{n}_3}$, however, are not generally unbiased and do not generally have minimum variance except for sufficiently large sample sizes.

Taking the case where $\boldsymbol{R}$ is a 3-D vector and σ_N is negligible, for example, the joint distribution for $\hat{\boldsymbol{x}}$ is

$$P_{\hat{\boldsymbol{x}}}\left(\hat{\boldsymbol{x}}\middle|\boldsymbol{x}\right) = \left|\boldsymbol{s}_1 \cdot \boldsymbol{s}_2 \times \boldsymbol{s}_3\right| P_{\boldsymbol{R}}\left(\mathbf{S}\hat{\boldsymbol{x}}\middle|\sigma\right), \tag{30}$$

which leads to the respective joint distributions

$$P_{grad}(\hat{p}_n, \hat{q}_n, \hat{\rho}|p_n, q_n, \rho) = \\ \frac{P_{\hat{\boldsymbol{x}}}(-\hat{\rho}\hat{p}_n(1+\hat{p}_n^2+\hat{q}_n^2)^{-1/2}, -\hat{\rho}\hat{q}_n(1+\hat{p}_n^2+\hat{q}_n^2)^{-1/2}, \hat{\rho}(1+\hat{p}_n^2+\hat{q}_n^2)^{-1/2}|\boldsymbol{x})}{\left|\hat{\rho}^2(1+\hat{p}_n^2+\hat{q}_n^2)^{-3/2}\right|}, \tag{31}$$

and

$$P_{polar}(\hat{\theta}_n, \hat{\phi}_n, \hat{\rho}|\theta_n, \phi_n, \rho) = \frac{P_{\hat{\boldsymbol{x}}}(\hat{\rho}\cos\hat{\phi}_n \sin\hat{\theta}_n, \hat{\rho}\sin\hat{\phi}_n \sin\hat{\theta}_n, \hat{\rho}\cos\hat{\theta}_n|\boldsymbol{x})}{\left|\hat{\rho}^2 \sin^2\hat{\theta}_n\right|}, \tag{32}$$

for the gradient and polar coordinate maximum likelihood estimates.

Returning to the general case when $\boldsymbol{R}$ is an N-D vector, the asymptotic maximum likelihood distributions for surface orientation and albedo follow (8) when (10) and (11) are satisfied.

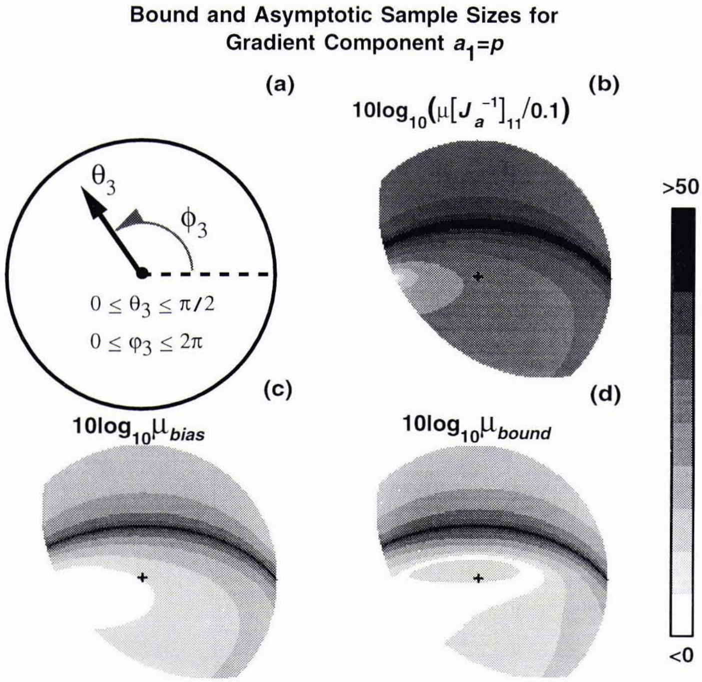

Figure 1. A visualization for $\boldsymbol{a}$ corresponding to the surface gradient parametrization for a 3-D measurement vector $\boldsymbol{R}$ with $\boldsymbol{\sigma}_N$ negligible and Lambertian surface defined by the polar coordinate parameterization ($\theta_n=\pi/4$, $\phi_n=\pi/4$, $\rho=1$). Incident vectors s_1 and s_2 are fixed by cone and polar angles ($\theta_1=\pi/2$, $\phi_1=0$), ($\theta_2=\pi/4$, $\phi_2=\pi/4$) but s_3 is allowed to vary as in (a) where the positive z-axis is central and points out of the page. (b) The bound $\left[J_a^{-1}\right]_1$ on the x-gradient $a_1=p$, including full 3-D coupling. Only values where $s_3 \cdot \boldsymbol{n}$ is positive and the Lambertian surface is in view from the positive z-axis are shown. Optimal resolution occurs when s_3 is tangent to the Lambertian surface along the x axis (horizontal), for the p-bound, and the y axis (vertical), for the q-bound (not shown), with sign so as to maximize the volume of incident vectors. Poorest resolution occurs when the volume of incident vectors approaches zero, as realized along the dark arc. (c) The number of samples necessary for the maximum likelihood estimate $\hat{a}_1 = \hat{p}$ to be effectively unbiased, from (10). (d) The number of samples necessary, from (11), for $\hat{p}$ to effectively attain the bound given in (b).

These computations exhibit the dramatic nonlinear variations in sample size necessary to first obtain an unbiased estimate and then a minimum variance estimate of Lambertian surface orientation and albedo from multi-static sonar images corrupted by signal-dependent speckle noise.

References

[1] N. C. Makris, "The effect of saturated transmission scintillation on ocean-acoustic intensity measurements," J. Acoust. Soc. Am. 100, 769-783, 1996.

[2] N. C. Makris, "Resolving Lambertian surface orientation from fluctuating radiance," J. Optical Soc. Am. (unpublished).

[3] N. C. Makris, "A foundation for logarithmic measures of fluctuating intensity in pattern recognition," Optics Letters 20, 2012-2014, 1995.

[4] R. A. Fisher, *Statistical Methods and Scientific Inference*, Hafner Publishing Co., New York, 1956.