Published in final edited form as:

Proc SPIE Int Soc Opt Eng. 2006 January 1; 6272: 62721W-. doi:10.1117/12.672798.

Task Performance in Astronomical Adaptive Optics

Harrison H. Barrett,

College of Optical Sciences and Department of Radiology, University of Arizona, Tucson AZ 85724

Kyle J. Myers,

NIBIB/CDRH Laboratory for the Assessment of Medical Imaging Systems, Food and Drug Administration, Rockville, MD 20850

Nicholas Devaney,

Department of Physics, National University of Ireland, Galway, Ireland

J. C. Dainty, and

Department of Physics, National University of Ireland, Galway, Ireland

Luca Caucci

Department of Electrical and Computer Engineering, University of Arizona, Tucson AZ 85721

Abstract

In objective or task-based assessment of image quality, figures of merit are defined by the performance of some specific observer on some task of scientific interest. This methodology is well established in medical imaging but is just beginning to be applied in astronomy. In this paper we survey the theory needed to understand the performance of ideal or ideal-linear (Hotelling) observers on detection tasks with adaptive-optical data. The theory is illustrated by discussing its application to detection of exoplanets from a sequence of short-exposure images.

Keywords

Adaptive optics; image quality; covariance; detection; exoplanets; Hotelling observer

1. INTRODUCTION

Adaptive optical (AO) systems have reached a high level of sophistication, and many ingenious system configurations and data-processing algorithms have been proposed. Methods for specifying system performance and image quality, on the other hand, seldom go beyond Strehl ratio. Indeed, it is often said that the goal of adaptive optics is to improve the Strehl ratio.

In this paper we take a different view, one adopted from the medical-imaging literature. In this view, the goal of a medical imaging system is to perform a specific task of clinical

interest, and the system performance is determined by how well the task can be performed. Similarly, astronomical AO systems are used for specific tasks of scientific interest, and it is reasonable to define image quality in terms of performance on these tasks. This approach has proven valuable for evaluating and optimizing radiological imaging systems,¹ and it is the premise of this paper that it will also be useful in adaptive optics.

In both fields, the tasks of interest are either *classification* or *estimation*. In a classification task, the goal is to assign the object that produced the image to one of two or more classes or hypotheses. In an estimation task, the goal is to determine numerical values of one or more parameters describing the object. The method by which the task is performed is called the *observer*.

An example of a classification task in medical imaging is detection of a tumor in a mammogram. The observer for this task can be either a human (a radiologist) or an automated detection algorithm, and in either case task performance can be specified by an ROC (receiver operating characteristic) curve, described briefly in Sec. 2.

An example of an estimation task in medical imaging is determination of the uptake of a targeted radiotracer that binds preferentially to a specific neuroreceptor. The data set for this task is a reconstructed tomographic image obtained by PET (positron emission tomography) or SPECT (single-photon emission computed tomography), and various mathematical observers are used to determine the uptake from the reconstructed image. Performance in this case can be defined either by accuracy of the estimate (including both bias and variance) or by the usefulness of the estimated value in performing a subsequent diagnostic classification.

Detection of a faint object such as an asteroid on a cluttered star field has much in common with detection of a small tumor in a mammogram, where normal anatomical structures obscure the tumor in the same way as background astronomical objects obscure the asteroid. Similarly, detection of an exoplanet is analogous to detection of a metastatic bone tumor on a rib in a SPECT image when the tumor is adjacent to the heart; like the host star in astronomy, the heart appears very bright since it is large and emits photons of the same energy as does the tumor, and it can be challenging to separate the two in a noisy, low-resolution image.

A common estimation task in astronomy is determination of the brightness of a star. Known as photometry, this task is completely analogous to assaying the local uptake of a radiotracer in SPECT or PET imaging.

For both estimation and classification tasks, calculation of task performance (in either medicine or astronomy) requires detailed knowledge of all statistical effects that influence the image or the observer. In radiology, the images are random because of photon noise from the discrete x rays or gamma rays, but also because the objects themselves are random. The photon noise in raw projection data is well described by independent Poisson statistics, but tomographic reconstruction algorithms lead to spatial correlations and more complicated image statistics. Object randomness includes normal anatomical structures as well as randomness in the tumor or other object of diagnostic interest, and these stochastic processes

must also be transformed through the image-forming elements and any post-processing such as tomographic reconstruction. The images are said to be doubly stochastic¹ since there is randomness from both object variability and measurement noise.

In AO systems, there is an additional source of randomness, the point-spread function (PSF) of the imaging system itself. If the adaptive optics functioned perfectly, this PSF would be a nonrandom Airy pattern, but residual speckle and unknown residual aberrations make it imperative to treat the PSF as random, so AO images are triply stochastic. Even though one of the three stochastic components, the Poisson noise from discrete photoelectric events, is statistically independent from measurement to measurement, the contributions from object and PSF randomness produce correlations and complicated statistical distributions.

In a paper recently submitted for publication,² we presented expressions for the mean data and spatiotemporal covariance matrices for general AO systems with all three sources of randomness. We then related these expressions to three specific tasks of interest in astronomical adaptive optics: (1) detection of a faint point object on a complex background with random spatial and temporal structure; (2) detection of a faint companion at an unknown location around a bright star, and (3) photometry of a given star in a crowded field. In all three cases, expressions were derived for the form of the optimal linear observer for performing the task as well as for the final performance achievable. Practical ways of calculating or estimating these covariance components were discussed, and ways of using them to compute task-based measures of image quality were presented.

The present paper concentrates on detection tasks, but there is a brief discussion of joint detection and estimation. It expands considerably on the analysis of Ref. 2 by looking at the Bayesian ideal observer, which almost always requires nonlinear operations on the data. The example of detection of an exoplanet from a sequence of short-exposure images is treated in some detail. Previous work on Bayesian detection of exoplanets^{3, 4} has considered only a single long-exposure image and a much simpler stochastic model.

Section 2 introduces the notation and some important definitions relating to triply stochastic spatiotemporal data, and Sec. 3 briefly reviews some key concepts from statistical decision theory. Section 4 summarizes results from ref. 2 on triply stochastic covariance matrices for AO systems and adds some discussion of probability density functions. Section 5 shows how these ideas can be applied to detection of an exoplanet.

2. DEFINITIONS AND NOTATION

This section introduces the notation that will be used in this paper, reviews the basic principles of multiply stochastic data, and discusses the nontrivial concept of linearity.

2.1. Doubly stochastic spatial data

Consider a digital imaging system viewing a time-independent object $f(\mathbf{r})$, where $\mathbf{r} = (x, y)$ is a 2D position vector. We shall often think of this object as a vector in a Hilbert space, in which case we shall denote it as $\mathbf{f}$.

A single digital image consists of M pixel values, $\{g_m, m = 1, \dots, M\}$, which can be considered as the components of an $M \times 1$ vector $\mathbf{g}$. As discussed above, this vector is random because of measurement noise and because different objects, selected at random from some ensemble of objects, will be imaged on different repetitions of the experiment. We can define the statistical expectation of this image with respect to both sources of randomness as

$$\bar{\mathbf{g}} \equiv \int d\mathbf{g} \, \text{pr}(\mathbf{g}) \mathbf{g}, \quad (2.1)$$

where $\text{pr}(\mathbf{g})$ is the overall M -dimensional probability density function (PDF) for the data, and the integral runs over all values that can be assumed by each of the components of $\mathbf{g}$.

We can indicate the doubly stochastic nature of the average by writing

$$\text{pr}(\mathbf{g}) = \int d\mathbf{f} \, \text{pr}(\mathbf{g}|\mathbf{f}) \, \text{pr}(\mathbf{f}), \quad (2.2)$$

where $\text{pr}(\mathbf{f})$ is an abstract way of indicating the PDF on all parameters needed to express the statistics of the random process $f(\mathbf{r})$. Thus,

$$\bar{\mathbf{g}} = \int d\mathbf{g} \int d\mathbf{f} \, \text{pr}(\mathbf{g}|\mathbf{f}) \, \text{pr}(\mathbf{f}) \mathbf{g} \equiv \langle \langle \mathbf{g} \rangle_{\mathbf{g}|\mathbf{f}} \rangle_{\mathbf{f}}. \quad (2.3)$$

We can also define a conditional average over the measurement noise alone by

$$\bar{\mathbf{g}}(\mathbf{f}) \equiv \langle \mathbf{g} \rangle_{\mathbf{g}|\mathbf{f}} = \int d\mathbf{g} \, \text{pr}(\mathbf{g}|\mathbf{f}) \mathbf{g}. \quad (2.4)$$

The imaging system is said to be linear if $\bar{\mathbf{g}}(\mathbf{f})$ is a linear function of $\mathbf{f}$, in which case the components of $\bar{\mathbf{g}}$ can be expressed as

$$\bar{g}_m = \int_{\infty} d^2r \, h_m(\mathbf{r}) f(\mathbf{r}), \quad (2.5)$$

where the kernel $h_m(\mathbf{r})$ is the response at pixel m for a point source at $\mathbf{r}$. The subscript ∞ indicates that the integral runs over the entire x - y plane.

The overall average image, in component form, is given by

$$\bar{\bar{g}}_m = \int_{\infty} d^2r \, h_m(\mathbf{r}) \bar{f}(\mathbf{r}), \quad (2.6)$$

where $\bar{f}(\mathbf{r})$ is the ensemble-average object.

2.2. Triply stochastic spatiotemporal data

Now consider an AO system that produces a temporal sequence of J images while viewing a time-dependent incoherent object $f(\mathbf{r}, t)$. The j^{th} image in the sequence is indicated by the $M \times 1$ vector $\mathbf{g}^{(j)}$, and the entire sequence is denoted $\mathbf{G}$.

By analogy to the doubly stochastic case, we shall use a single overbar to indicate an average over the measurement noise alone, conditional on a specific object and PSF. For incoherent imaging, this average is a linear function of the spatiotemporal object, given by

$$\bar{g}_m^{(j)} = \int d^2 r_d d_m(\mathbf{r}_d) \int_{t_j}^{t_j+T} dt \int d^2 r p(\mathbf{r}_d, \mathbf{r}, t) f(\mathbf{r}, t), \quad (2.7)$$

where $\mathbf{r}_d$ is a 2D vector in the image plane, T is the frame time, $d_m(\mathbf{r}_d)$ describes the sensitivity of the m^{th} detector pixel, and $p(\mathbf{r}_d, \mathbf{r}, t)$ is the random, anisoplanatic, time-dependent, incoherent PSF of the system.

As in ref. 2, we shall assume that the object is a slowly varying function of time, essentially constant over one frame of the science camera, in which case (2.6) becomes

$$\bar{g}_m^{(j)} = \int_\infty d^2 r h_m^{(j)}(\mathbf{r}) f^{(j)}(\mathbf{r}), \quad m=1, \dots, M, j=1, \dots, J, \quad (2.8)$$

where $f^{(j)}(\mathbf{r}) = f(\mathbf{r}, t_j)$,

$$h_m^{(j)}(\mathbf{r}) = \int d^2 r_d d_m(\mathbf{r}_d) p^{(j)}(\mathbf{r}_d, \mathbf{r}), \quad (2.9)$$

and

$$p^{(j)}(\mathbf{r}_d, \mathbf{r}) = \int_{t_j}^{t_j+T} dt p(\mathbf{r}_d, \mathbf{r}, t). \quad (2.10)$$

The sequence of object functions, $\{f^{(j)}(\mathbf{r}), j=1, \dots, J\}$, will be denoted $\mathbf{F}$, and the sequence of PSFs, $\{p^{(j)}(\mathbf{r}_d, \mathbf{r}), j=1, \dots, J\}$, will be denoted $\mathbf{P}$.

A second overbar will be used to indicate an average over the random PSFs $\mathbf{P}$ conditional on $\mathbf{F}$. In component form,

$$\bar{\bar{g}}_m^{(j)} = \bar{\bar{g}}_m^{(j)}(\mathbf{f}^{(j)}) = \int d^2 r \bar{h}_m^{(j)}(\mathbf{r}) f^{(j)}(\mathbf{r}), \quad (2.11)$$

where the average kernel is related to the average incoherent PSF by

$$\bar{h}_m^{(j)}(\mathbf{r}) = \int d^2 r_d d_m(\mathbf{r}_d) \bar{p}^{(j)}(\mathbf{r}_d, \mathbf{r}). \quad (2.12)$$

The final average, over the object variability, yields

$$\bar{\bar{\bar{g}}}_m^{(j)} = \int d^2 r \left\langle \bar{h}_m^{(j)}(\mathbf{r}) f^{(j)}(\mathbf{r}) \right\rangle_{\mathbf{F}} = \int d^2 r \bar{h}_m^{(j)}(\mathbf{r}) \bar{f}^{(j)}(\mathbf{r}), \quad (2.13)$$

where the second form holds only if $\mathbf{p}^{(j)}$ is statistically independent of $\mathbf{F}$.

Some special cases of these averages will be discussed in Sec. 4.

3. CONCEPTS FROM STATISTICAL DECISION THEORY

In this section we briefly review some basic concepts from statistical decision theory, relating them specifically to triply stochastic spatiotemporal data.

3.1. Detection problems and ROC curves

In any classification task, the goal is to assign the object that produced an image to one of two or more classes. If the hypothesis that the object belongs to the k^{th} class is denoted H_k , then the probability law for the data when hypothesis H_k is true is denoted $\text{pr}(\mathbf{G}|H_k)$. In a signal-detection task, the hypotheses are signal-absent and signal-present.

If we assume that each image must be assigned without equivocation either to hypothesis H_0 (signal absent) or to H_1 (signal present), the decision on a detection task can be made in complete generality by computing some scalar test statistic $t(\mathbf{G})$ from the data; the observer then decides H_1 if the test statistic is greater than a decision threshold and decides H_0 otherwise. The value of the threshold controls the tradeoff between true positive decisions (correctly choosing H_1) and false positive decisions (choosing H_1 when H_0 is true). In signal-detection problems, the true-positive fraction (TPF) is called the *probability of detection*, and the false-positive fraction (FPF) is called the *false-alarm rate*.

A plot of TPF vs. FPF as the threshold is varied is called a *receiver operating characteristic* (ROC) curve. Meaningful figures of merit for binary classification include the true positive fraction at a specified false-positive fraction (the Neyman-Pearson criterion), the area under the ROC curve (AUC), and certain detectability indices derived from the ROC curve. The probability of detection alone is not a meaningful metric since it can always be made large, even unity, simply by choosing a low threshold.

3.2. Random signals, backgrounds and PSFs

In the medical imaging literature, it is common to divide detection tasks further depending on whether the background object, the signal to be detected or both are random. The most tractable – and least realistic – problem is where the background is specified exactly and the signal, if present, is also specified exactly; such problems are known as SKE/BKE (signal known exactly, background known exactly). The only remaining randomness in an SKE/BKE problem is the measurement noise and, of course, whether or not the signal is present. Though SKE/BKE problems are never realistic models for the actual clinical tasks of interest, we shall see that they are necessary stepping stones to the analysis of more practical problems.

In clinical practice, the background is not known. In mammography, for example, the normal anatomical structures vary randomly from patient to patient and must be treated as a spatial random process. To make any progress on this problem, we must have at least partial knowledge of the statistics of this random process, so in this case we say that the problem is BKS (background known statistically).

Similarly, in real clinical applications, the signal to be detected (e.g., a possible tumor) is also random in size, shape and location. Again, we must have some partial statistical

knowledge of these parameter in order to even be able to define the signal part of the image, so this situation is referred to as SKS (signal known statistically).

These distinctions are needed in astronomy as well. In planet detection, for example, the planet brightness and location are not known ahead of time, so the problem is inherently SKS. Sky background and other astronomical objects contribute to a random background, so real astronomical problems are also BKS.

In astronomy, however, there is the additional complication of the random PSF. In a perfect AO system or in a long-exposure image where the PSF fluctuations average out, the PSF is nonrandom and the problem is designated PKE (PSF known exactly). On the other hand, if the residual speckle is important as in a sequence of short-exposure images, the problem is PKS (PSF known statistically).

3.3. Ideal Bayesian observer

The ideal observer on a detection task is defined variously as one that maximizes the area under the ROC curve, maximizes the TPF at all specified FPFs or minimizes a cost function defined in terms of TPF and FPF. By any of these criteria, the test statistic used by the ideal observer is the likelihood ratio,

$$\Lambda(\mathbf{G}) \equiv \frac{\text{pr}(\mathbf{G}|H_1)}{\text{pr}(\mathbf{G}|H_0)}. \quad (3.1)$$

Equivalently, the ideal observer can use the logarithm of the likelihood ratio, $\lambda(\mathbf{G}) \equiv \ln \Lambda(\mathbf{G})$, as the test statistic. In either case, the ideal observer computes the test statistic and compares it to a threshold, deciding that H_1 is true if the statistic exceeds the threshold. Some optimization criteria fix the value of the threshold according to the prevalences of the two hypotheses or costs assigned to erroneous decisions, but in every case the ideal test statistic is the likelihood ratio or its logarithm.

In practice, the relevant probabilities $\text{pr}(\mathbf{G}|H_k)$ in (3.1) are seldom known directly. Instead, they are given by

$$\text{pr}(\mathbf{G}|H_k) = \int d\mathbf{F} \text{pr}(\mathbf{G}|\mathbf{F}) \text{pr}(\mathbf{F}|H_k) = \langle \text{pr}(\mathbf{G}|\mathbf{F}) \rangle_{\mathbf{F}|H_k}, \quad (k=0,1). \quad (3.2)$$

For a PKE problem, $\text{pr}(\mathbf{G}|\mathbf{F})$ describes the measurement noise, and it is easy to compute from the basic physics of the problem. For detectors with no significant readout noise, for example, $\text{pr}(\mathbf{G}|\mathbf{F})$ describes independent Poisson random variables at each pixel in each frame. With no PSF randomness and the specific object $\mathbf{F}$, knowledge of the mean image at each pixel and frame is sufficient to determine $\text{pr}(\mathbf{G}|\mathbf{F})$.

If the PSF is random, it is useful to write

$$\text{pr}(\mathbf{G}|H_k) = \int d\mathbf{P} \int d\mathbf{F} \text{pr}(\mathbf{G}|\mathbf{P}, \mathbf{F}) \text{pr}(\mathbf{P}|\mathbf{F}) \text{pr}(\mathbf{F}|H_k) = \langle \text{pr}(\mathbf{G}|\mathbf{P}, \mathbf{F}) \rangle_{\mathbf{P}|\mathbf{F}} \bigg|_{\mathbf{F}|H_k}. \quad (3.3)$$

The conditioning of the average over PSF $\mathbf{P}$ on the object $\mathbf{F}$ is needed in an AO system if the control signals for the adaptive element are derived from the object of interest; if a separate guide star is used and not treated as a part of $\mathbf{F}$, then the object and PSF are statistically independent. In either case, $\text{pr}(\mathbf{G}|\mathbf{P}, \mathbf{F})$ is easily derived from the physics of the problem. Markov-chain Monte Carlo (MCMC) methods can in many cases be used to perform the two averages.¹

3.4. Ideal linear observer

Often the likelihood ratio is difficult to compute, and in these cases a useful alternative to the ideal observer is the ideal linear observer, called the Hotelling observer^{1, 5–7} in the medical literature.

Linear observers compute linear discriminants, so the test statistic for an image sequence has the form $t(\mathbf{G}) = \mathbf{W}'\mathbf{G}$, where $\mathbf{W}$ is another image sequence called the template. The notation $\mathbf{W}'\mathbf{G}$ denotes a pixel-by-pixel, frame-by-frame scalar product of the template with the spatiotemporal data.

The Hotelling discriminant uses a template that maximizes a certain class separability measure.⁸ Linear test statistics are usually normally distributed by dint of the central limit theorem, and in this case maximizing this class separability is equivalent to maximizing the AUC among linear observers.

It can also be shown that the Hotelling test statistic is equal to the log-likelihood ratio if the raw data are normal with the same covariance under both hypotheses, so the Hotelling observer is identical to the ideal observer in this case and thus maximizes AUC among all observers, not just linear ones.

Computation of the Hotelling test statistic requires only the overall mean vectors and the covariance matrices of the data under the two hypotheses. The test statistic is given by

$$t_{\text{Hot}}(\mathbf{G}) = \mathbf{W}'\mathbf{G} = [\bar{\mathbf{G}}_1 - \bar{\mathbf{G}}_0]^t \mathbf{K}_{av}^{-1} \mathbf{G}, \quad \text{where} \quad \mathbf{K}_{av} \equiv \frac{1}{2} [\mathbf{K}_{\mathbf{G}|H_1} + \mathbf{K}_{\mathbf{G}|H_0}]. \quad (3.4)$$

The inverse of the average covariance matrix is related to the familiar signal-processing operation of prewhitening, and for this reason, the Hotelling observer is sometimes called a prewhitening matched filter; unless the noise is stationary, however, the prewhitening and matched filtering cannot be carried out in the Fourier domain.

A figure of merit for the Hotelling observer is the Hotelling signal-to-noise ratio, sometimes called the Hotelling trace; it is given by

$$\text{SNR}_{\text{Hot}}^2 \equiv [\bar{\mathbf{G}}_1 - \bar{\mathbf{G}}_0]^t \mathbf{K}_{av}^{-1} [\bar{\mathbf{G}}_1 - \bar{\mathbf{G}}_0] = \text{tr} \left\{ \mathbf{K}_{av}^{-1} [\bar{\mathbf{G}}_1 - \bar{\mathbf{G}}_0] [\bar{\mathbf{G}}_1 - \bar{\mathbf{G}}_0]^t \right\}, \quad (3.5)$$

where $\text{tr}\{\cdot\}$ denotes the trace (sum of the diagonal elements) of the matrix.

Computational approaches to evaluating the Hotelling test statistic and SNR are discussed in refs. 1 and 2.

3.5. Detection with location uncertainty

In most detection problems, we want to know not only that a signal is present but where it is. A useful alternative to the ROC curve in these cases is the LROC (localization ROC) curve, which plots the probability of detection and correct localization (within some preset tolerance) against the false-alarm rate. In the medical literature, the common figure of merit for joint detection/localization is the area under the LROC curve (ALROC).

As developed so far, neither the ideal nor the Hotelling observer gives any information about signal location, but both are readily modified to do so. A well-known way of using the likelihood for any detection problem with parameter uncertainty is the generalized likelihood ratio test (GLRT), in which the numerator in (3.1) is modified to

$$\Lambda(\mathbf{G}; \boldsymbol{\theta}) \equiv \frac{\text{pr}(\mathbf{G} | \boldsymbol{\theta}, H_1)}{\text{pr}(\mathbf{G} | H_0)}, \quad (3.6)$$

where the vector $\boldsymbol{\theta}$ specifies the location or any other random parameter of the signal. In a GLRT, the parameter-dependent likelihood ratio $\Lambda(\mathbf{G}; \boldsymbol{\theta})$ is first maximized with respect to $\boldsymbol{\theta}$, and the maximized result is compared to a threshold to decide if the signal is present; if that decision is positive, the maximum value of $\boldsymbol{\theta}$ is chosen as the estimate of the signal parameters.

For the case where $\boldsymbol{\theta}$ is signal location, Khurd and Gindi⁹ have shown that a modified GLRT, where the likelihood ratio is weighted by the probability of occurrence of each location, maximizes ALROC.

The Hotelling observer alone is ineffective with location uncertainty because the average signal, $\bar{\mathbf{G}}_1 - \bar{\mathbf{G}}_0$ in (3.4) or (3.5), tends to a small constant, independent of location in the image, when the double-bar average includes random signal location. As with the GLRT, however, a simple modification salvages the concept. The Hotelling template can be defined for an SKE problem with a specified location. The template can then be scanned over all possible locations, and the resulting maximum value becomes the final test statistic for the detection decision. This strategy, called the scanning Hotelling observer, is discussed in more detail in refs. 1 and 2.

4. STATISTICAL PROPERTIES OF AO IMAGES

As we saw in the previous section, the covariance matrices under both hypotheses are needed for the Hotelling or scanning Hotelling observer, and the full PDFs of the data are

needed for the ideal observer or the GLRT. In this section we discuss some approaches to obtaining these statistical descriptions.

4.1. Covariance matrix

The overall covariance matrix of a triply stochastic image sequence is defined as

$$\mathbf{K}_G \equiv \left\langle \left[\mathbf{G} - \overline{\overline{\mathbf{G}}} \right] \left[\mathbf{G} - \overline{\overline{\mathbf{G}}} \right]^t \right\rangle_{\mathbf{G}|\mathbf{P},\mathbf{F}} \quad n = \left\langle \left\langle \left\langle \left[\mathbf{G} - \overline{\overline{\mathbf{G}}} \right] \left[\mathbf{G} - \overline{\overline{\mathbf{G}}} \right]^t \right\rangle_{\mathbf{G}|\mathbf{P},\mathbf{F}} \right\rangle_{\mathbf{P}|\mathbf{F}} \right\rangle_{\mathbf{F}}. \quad (4.1)$$

To be explicit, $\mathbf{K}_G$ is an $MJ \times MJ$ matrix with components given by

$$[\mathbf{K}_G]_{mm'}^{(j,j')} = \left\langle \left\langle \left\langle \left[g_m^{(j)} - \overline{\overline{g}}_m^{(j)} \right] \left[g_{m'}^{(j')} - \overline{\overline{g}}_{m'}^{(j')} \right] \right\rangle_{\mathbf{G}|\mathbf{P},\mathbf{F}} \right\rangle_{\mathbf{P}|\mathbf{F}} \right\rangle_{\mathbf{F}}. \quad (4.2)$$

Adding and subtracting terms in each factor of (4.2), we obtain

$$\mathbf{K}_G = \left\langle \left\langle \left\langle \left[\mathbf{G} - \overline{\mathbf{G}} + \overline{\mathbf{G}} - \overline{\overline{\mathbf{G}}} + \overline{\overline{\mathbf{G}}} - \overline{\overline{\mathbf{G}}} \right] \left[\mathbf{G} - \overline{\mathbf{G}} + \overline{\mathbf{G}} - \overline{\overline{\mathbf{G}}} + \overline{\overline{\mathbf{G}}} - \overline{\overline{\mathbf{G}}} \right]^t \right\rangle_{\mathbf{G}|\mathbf{P},\mathbf{F}} \right\rangle_{\mathbf{P}|\mathbf{F}} \right\rangle_{\mathbf{F}}. \quad (4.3)$$

Even without any assumptions of independence, the cross terms vanish identically, and we can write

$$\mathbf{K}_G = \overline{\overline{\mathbf{K}}}_G^{noise} + \overline{\overline{\mathbf{K}}}_G^{PSF} + \mathbf{K}_{\overline{\overline{\mathbf{G}}}}^{obj}, \quad (4.4)$$

where

$$\overline{\overline{\mathbf{K}}}_G^{noise} \equiv \left\langle \left\langle \left\langle \left[\mathbf{G} - \overline{\mathbf{G}} \right] \left[\mathbf{G} - \overline{\mathbf{G}} \right]^t \right\rangle_{\mathbf{G}|\mathbf{P},\mathbf{F}} \right\rangle_{\mathbf{P}|\mathbf{F}} \right\rangle_{\mathbf{F}}; \quad (4.5)$$

$$\overline{\overline{\mathbf{K}}}_G^{PSF} \equiv \left\langle \left\langle \left[\overline{\mathbf{G}} - \overline{\overline{\mathbf{G}}} \right] \left[\overline{\mathbf{G}} - \overline{\overline{\mathbf{G}}} \right]^t \right\rangle_{\mathbf{P}|\mathbf{F}} \right\rangle_{\mathbf{F}}; \quad (4.6)$$

$$\mathbf{K}_{\overline{\overline{\mathbf{G}}}}^{obj} \equiv \left\langle \left[\overline{\overline{\mathbf{G}}} - \overline{\overline{\overline{\mathbf{G}}}} \right] \left[\overline{\overline{\mathbf{G}}} - \overline{\overline{\overline{\mathbf{G}}}} \right]^t \right\rangle_{\mathbf{F}}. \quad (4.7)$$

Thus the overall covariance matrix for a triply stochastic image sequence can be rigorously decomposed into three terms representing, respectively, the contributions from measurement noise, from the random PSF and from randomness in the object being imaged.

The noise term in the covariance is almost always diagonal: measurements at different detector pixels and/or different frames are uncorrelated. The other two terms, however, exhibit significant spatial and temporal correlations; approaches to estimating them by simulation are discussed in ref. 2.

4.2. Probability density functions

Multivariate normal random processes are fully specified by their mean vector and covariance matrices, but normal statistics cannot accurately describe incoherent images since irradiance cannot go negative. Thus, if we wish to compute the performance of an ideal Bayesian observer on a detection task, we need more information than just the mean and covariance. This section briefly discusses some approaches to obtaining relevant probability density functions.

4.2.1. Background models—The simplest model for a random sky background is that it is spatially constant over the field of view and does not vary with time over the observation period, but that the absolute level of the background is not known *a priori*. In that case, the background is described by a single scalar random variable f_b . The ideal observer needs information on the PDF $\text{pr}(f_b)$, and the Hotelling observer needs its mean and variance.

The next level of sophistication is to allow the spatially uniform background level to vary with time, denoting it $f_b(t)$. In many cases, this random process might be treated as temporally stationary. Its point statistics could be established experimentally for a particular observing situation.

Various background models with random spatial variation are used in medical imaging, and many of these yield images similar to astronomical scenes. In many cases, the models are not only realistic but also mathematically tractable; analytic expressions can be given for the characteristic functional, which in principle contains all statistical information about the random process. For a review of this topic, see Barrett and Myers.¹

One particular spatial model that should accurately model globular clusters and other aggregations of point-like object is a nonstationary Poisson point process.^{2, 3}

4.2.2. PSF models—How one models a random PSF depends on the exposure time T relative to the speckle correlation time of the PSF, denoted τ_c , and on how many images are included in the data set.

If the data consist of a single very long exposure, the PSF may not be known accurately, but it is fixed for that exposure and does not have to be treated as random. It can be estimated from measured data¹⁰ or calculated from knowledge of the atmosphere, telescope and AO system, and this estimated PSF can then be used for performing the tasks of interest.

If the whole observing period is divided into shorter exposures but each individual frame is still long compared to the speckle correlation time, then the PSF for one frame can be considered as random but with a form parameterized by the local atmospheric conditions during that exposure. It may even be possible to specify the form of the PSF adequately by

knowing the average Fried parameter r_0 for that frame. In that case the sequence $\mathbf{P}$ is temporally correlated but fully specified by the instantaneous Fried parameter for each frame. The problem of modeling the PSF boils down to understanding the statistics of the scalar random process $r_0(t)$.

The more interesting problem is when the data $\mathbf{G}$ consist of many short exposures, each with $T < \tau_c$, so that $\mathbf{P}$ is a full spatiotemporal random process. More details on this case are given in Sec. 5.

5. EXAMPLE: DETECTION OF EXOPLANETS

One way to formulate an interesting planet-detection problem is to assume that the astronomical scene of interest consists of a uniform nonrandom background, a single star, and possibly a single planet, and that the data are a sequence of short-exposure images. The randomness in the data then comes from measurement noise, the residual speckle in the PSF and the random size and location of the planet. The problem can be classified as SKS/BKE/PSK in the language of Sec. 3.2. We shall use this problem to illustrate the statistical approaches described above.

Numerous authors have analyzed the deterministic and statistical properties of residual speckle in AO systems^{11–14} None of these authors, however, consider either the full multivariate PSF needed by the ideal observer or even the spatiotemporal covariance needed by the Hotelling observer.

5.1. Pinned speckle

For simplicity, assume that the star is located at $\mathbf{r} = \mathbf{0}$, and assume that the light is narrowband with mean wavelength λ . If the instantaneous pupil function, after the AO correction, is given by $a_{ap}(\mathbf{r}) \exp[i\phi(\mathbf{r}, t)]$, where $a_{ap}(\mathbf{r})$ is the aperture function of the telescope and $\phi(\mathbf{r}, t)$ is the residual uncorrected phase, then the complex scalar field at point $\mathbf{r}_d$ in the image (or detector) plane is given, within the Fresnel approximation, by

$$u(\mathbf{r}_d, t) \propto \int_{\infty} d^2r a_{ap}(\mathbf{r}) \exp[i\phi(\mathbf{r}, t)] \exp\left(\frac{2\pi i \mathbf{r} \cdot \mathbf{r}_d}{\lambda f}\right). \quad (5.1)$$

By expanding $\exp[i\phi(\mathbf{r}, t)]$ in a Taylor series, taking the squared modulus of the field, and retaining terms linear or quadratic in the residual phase, we can express the random irradiance in the image plane approximately as¹³

$$I(\mathbf{r}_d, t) \propto |A_{ap}(\tilde{\mathbf{r}})|^2 - 2A_{ap}(\tilde{\mathbf{r}}) \text{Im} \{ [A_{ap} * \Phi](\tilde{\mathbf{r}}, t) \} + |[A_{ap} * \Phi](\tilde{\mathbf{r}}, t)|^2 - A_{ap}(\tilde{\mathbf{r}}) \text{Re} \{ [A_{ap} * \Phi * \Phi](\tilde{\mathbf{r}}, t) \}, \quad (5.2)$$

where $A_{ap}(\boldsymbol{\rho}, t)$ and $\Phi(\boldsymbol{\rho}, t)$ are the 2D spatial Fourier transforms of $a(\mathbf{r})$ and $\phi(\mathbf{r}, t)$, respectively, and $\tilde{\mathbf{r}} \equiv \mathbf{r}_d/(\lambda f)$. Since there is only a single on-axis point source in the field, this expression is the random PSF, $p(\mathbf{r}_d, \mathbf{r}, t)$ that appears in (2.7), but evaluated at $\mathbf{r} = \mathbf{0}$.

The leading term in (5.2) is the ideal (nonrandom) Airy pattern for no aberrations (which is a real function for a symmetric aperture), and the other terms, corresponding to residual speckle, are spatiotemporal random processes. The second and fourth terms are referred to as pinned speckle since they are modulated by the Airy pattern; the third term is not pinned since it is convolved with the Airy pattern but not multiplied by it. For later reference, we label the terms in (5.2), in sequence, as I_1 , I_2 , I_3 and I_4 . Thus $I = I_1 + I_2 + I_3 + I_4$.

5.2. Potentially useful assumptions

Next we look at three assumptions that will simplify computation of the spatiotemporal covariance and PDF of the data. All three derive from properties of uncorrected Kolmogorov turbulence, but it is argued that the same properties hold to a reasonable approximation after AO correction.

5.2.1. Spatial stationarity in the pupil—According to Kolmogorov theory, the uncorrected pupil phase fluctuations (before correction by the deformable mirror and truncation by the aperture) are well described as a wide-sense-stationary spatial random process, but it is not so obvious that the wave emerging from the deformable mirror is spatially stationary. The fixed locations of the actuators on the mirror argue against any claim that the spatial autocorrelation function of the field or the phase in the pupil depends only on the separation of the two observation points. A well-functioning AO system, however, will accurately correct all phase fluctuations within a subspace spanned by the deformable mirror modes. If the original random process is stationary and the component of the atmospheric phase in the mirror subspace is uncorrelated with the component in its orthogonal complement, then the corrected phase is stationary as well.

5.2.2. Frozen flow—The spatiotemporal statistics can be reduced to a purely spatial problem if we make the familiar Taylor frozen-flow assumption, under which $\phi(\mathbf{r}, t) = \phi_0(\mathbf{r} - \mathbf{v}t)$ and

$$\Phi(\boldsymbol{\rho}, t) = \Phi_0(\boldsymbol{\rho}) \exp(2\pi i \boldsymbol{\rho} \cdot \mathbf{v}t), \quad (5.3)$$

where $\mathbf{v}$ is the wind velocity. Normally the frozen-flow hypothesis is applied to uncorrected phase aberrations, but if the AO system responds rapidly and nearly nulls out the components in the subspace spanned by the mirror modes, the hypothesis should apply to the residual phase after correction as well.

5.2.3. Circular Gaussian—Another useful assumption is that $\Phi(\mathbf{r}, t)$ is a circular Gaussian spatial random process for a fixed time t . One justification is similar to the argument used above: The uncorrected phase is a Gaussian in Kolmogorov turbulence theory, and a good AO system projects the Kolmogorov phase onto the orthogonal complement of the subspace spanned by the mirror modes. That projection is a linear operator, and any linear transformation of a Gaussian is a Gaussian. Moreover, we can argue that the real and imaginary parts of $\Phi(\mathbf{r}, t)$ are identically distributed since a small shift of the uncompensated part of the pupil phase will interchange real and imaginary parts.

5.3. PSF term in the covariance matrix

Extensive simulations are underway to understand fully the PSF term in the covariance for the special case of pinned speckle and short-exposure images. Numerical results will be reported separately, but in this section we examine the problem analytically with the help of the assumptions discussed above. Particular attention will be paid to the effect of the linear term, I_2 , in (5.2).

5.3.1. Autocovariance of the irradiance—The spatiotemporal autocovariance function of the irradiance produced by a single point object on the optical axis is defined by

$$k_I(\mathbf{r}_d, \mathbf{r}'_d, t, t') \equiv \langle I(\mathbf{r}_d, t) I(\mathbf{r}'_d, t') \rangle - \langle I(\mathbf{r}_d, t) \rangle \langle I(\mathbf{r}'_d, t') \rangle. \quad (5.4)$$

Decomposing $I(\mathbf{r}_d, t)$ as in (5.2), using the fact that I_1 is nonrandom, and assuming that all odd moments of Φ are zero, we obtain

$$k_I = k_{22} + k_{33} + k_{44} + k_{34} + k_{43}, \quad (5.5)$$

where

$$k_{mn} \equiv \langle I_m(\mathbf{r}_d, t) I_n(\mathbf{r}'_d, t') \rangle - \langle I_m(\mathbf{r}_d, t) \rangle \langle I_n(\mathbf{r}'_d, t') \rangle. \quad (5.6)$$

All of these terms can be evaluated analytically if we are willing to make some of the assumptions introduced in Sec. 5.2. We illustrate this point by considering specifically k_{22} , given in the scaled coordinates ($\tilde{\mathbf{r}} = \mathbf{r}_d / \lambda f$) by

$$k_{22}(\tilde{\mathbf{r}}, t, \tilde{\mathbf{r}}', t') = 4A_{ap}(\tilde{\mathbf{r}})A_{ap}(\tilde{\mathbf{r}}') \langle \text{Im} \{ [A_{ap} * \Phi](\tilde{\mathbf{r}}, t) \} \text{Im} \{ [A_{ap} * \Phi](\tilde{\mathbf{r}}', t') \} \rangle. \quad (5.7)$$

In the frozen-flow approximation, the essence of the calculation is to compute expectations of the form,

$$\langle [A_{ap} * \Phi](\tilde{\mathbf{r}}, t) [A_{ap} * \Phi]^*(\tilde{\mathbf{r}}', t') \rangle = \int d^2 \rho \int d^2 \rho' A_{ap}(\tilde{\mathbf{r}} - \rho) A_{ap}(\tilde{\mathbf{r}}' - \rho') \exp[2\pi i(\rho \cdot \mathbf{v}t - \rho' \cdot \mathbf{v}t')] \langle \Phi_0(\rho) \Phi_0^*(\rho') \rangle. \quad (5.8)$$

If we also assume that $\phi_0(\mathbf{r})$ is wide-sense stationary, it can be shown that

$$\langle \Phi_0(\rho) \Phi_0^*(\rho') \rangle = S_{\phi_0}(\rho) \delta(\rho - \rho'), \quad (5.9)$$

where $S_{\phi_0}(\rho)$ is the power spectral density of $\phi_0(\mathbf{r})$. Then (5.8) becomes

$$\langle [A_{ap} * \Phi](\tilde{\mathbf{r}}, t) [A_{ap} * \Phi]^*(\tilde{\mathbf{r}}', t') \rangle = \int d^2 \rho A_{ap}(\tilde{\mathbf{r}} - \rho) A_{ap}(\tilde{\mathbf{r}}' - \rho) \exp[2\pi i \rho \cdot \mathbf{v}(t - t')] S_{\phi_0}(\rho). \quad (5.10)$$

The complex conjugate of (5.10) also appears in k_{22} , as do two terms involving $\langle \Phi_0(\rho)\Phi_0(\rho^*) \rangle$ and its conjugate. We can still use (5.9) for these latter expectations because $\Phi_0(\rho^*) = \Phi_0^*(-\rho)$, and the final answer is

$$k_{22}(\tilde{\mathbf{r}}, t, \tilde{\mathbf{r}}', t') = A_{ap}(\tilde{\mathbf{r}})A_{ap}(\tilde{\mathbf{r}}') \left\{ \int d^2\rho A_{ap}(\tilde{\mathbf{r}} - \rho)A_{ap}(\tilde{\mathbf{r}}' - \rho) \cos[2\pi\rho \cdot \mathbf{v}(t - t')] S_{\varphi_0}(\rho) - \int d^2\rho A_{ap}(\tilde{\mathbf{r}} - \rho)A_{ap}(\tilde{\mathbf{r}}' + \rho) \cos[2\pi\rho \cdot \mathbf{v}(t - t')] \right\} \quad (5.11)$$

This result predicts a stationary time dependence for k_{22} , and it also predicts a strong positive correlation for $t = t'$ and $\tilde{\mathbf{r}}'$ near $\tilde{\mathbf{r}}$ and a strong negative correlation for $\tilde{\mathbf{r}}'$ near $-\tilde{\mathbf{r}}$. Simulation studies to check these predictions are in progress.

The quadratic terms in (5.2) are more complicated, but their means and covariances are calculable from properties of circular Gaussians if we are willing to add that assumption.

5.4. Complete covariance matrix

To go from the autocovariance *function* of the PSF, $k_I(\mathbf{r}, \mathbf{r}', t, t')$, to the PSF term in the covariance *matrix*, we must integrate over frames and pixels as in (2.10) and (2.12).

In principle, we should also include an object variability term, but we have assumed in this section that the background and the host star are nonrandom, so there is no object variability under the planet-absent hypothesis, H_0 . There is variability in planet location and brightness under H_1 , but these factors come into the mean signal, not the average covariance, if the star is much brighter than its companion. Moreover, with the scanning Hotelling observer, the mean and covariance are conditioned on a particular location.

We do need to include a covariance term accounting for measurement noise. If the noise is statistically independent (conditional on the object and PSF) from pixel to pixel and frame to frame, we can write

$$\left[\overline{\mathbf{K}}_{\mathbf{G}}^{\text{noise}} \right]_{mm'}^{(j,j')} = \sigma_{mj}^2 \delta_{mm'} \delta_{jj'}. \quad (5.12)$$

If both readout and Poisson noise are present, then

$$\sigma_{mj}^2 = \sigma_m^2 + \overline{g}_m^{(j)}, \quad (5.13)$$

where σ_m^2 is the electronic noise for the m^{th} detector pixel in photoelectron units and $\overline{g}_m^{(j)}$ in the present problem is obtained by averaging (5.2) over random PSFs and normalizing to photoelectron units.

Thus, with only minimal numerical calculation, we will have access to the full covariance matrix needed to implement the Hotelling or scanning Hotelling observer.

5.5. Probability density functions

If we wish to go beyond Hotelling to the ideal observer, we need PDFs for the data, not just means and covariance. We have already suggested that it might be valid take Φ as circular Gaussian; in that case the I_2 term in (5.2) is fully characterized, as is its contribution to the PDF of measured data. It is still necessary, however, to validate the circular-Gaussian assumption with detailed simulations.

The statistics of the I_3 and I_4 terms are not so easy to determine, even under a circular Gaussian model, because they involve products of spatiotemporal random processes. Further theoretical investigation is needed if these terms are significant.

The probability law for the measurement noise (again conditional on an object and PSF) is easy to state in the limit of pure Poisson noise or pure Gaussian readout noise.

6. SUMMARY AND CONCLUSIONS

The main point stressed in this paper is that task performance is the ultimate measure of image quality. Astronomical images are produced either for estimation tasks, such as photometry, or for classification tasks such as planet detection. Both kinds of task were analyzed previously for AO systems in ref. 2, but the analysis was restricted to linear observers, namely the Hotelling observer for binary detection tasks and the generalized Wiener estimator for estimation tasks.

In this paper we applied the linear analysis of ref. 2 specifically to detection of exoplanets, and we extended the discussion of detection tasks to nonlinear ideal observers that maximize the area under the ROC curve. We also briefly discussed the LROC curve for joint detection and estimation tasks.

As in ref. 2, we noted that task performance in AO systems is limited by three main sources of data randomness: measurement noise, object randomness and residual PSF randomness. For short-exposure images, the data from an AO system is a discrete spatiotemporal random process. Calculation of achievable performance by any observer depends on accurate multivariate statistical knowledge. For linear observers, mean vectors and spatiotemporal covariance matrices are needed, and for the ideal observer full multivariate PDFs are required. Univariate statistics say very little about task performance. In particular, pixel signal-to-noise ratio is almost unrelated to detectability.

For the specific problem of exoplanet detection from a sequence of short-exposure images, we found that the relevant covariance matrices could be expressed in a numerically tractable form with the help of two plausible assumptions, namely frozen flow and spatial stationarity of the phase in the pupil. The range of validity of these assumptions remains to be explored.

As a result of the investigations reported here, we now have the ingredients needed to implement the Hotelling and scanning Hotelling observer for exoplanet detection. This future work should lead to improved detectability, but perhaps more importantly it will provide a firm basis for optimization of AO systems for this task. In particular, it will

provide a direct comparison of the achievable detectability with multiple short-exposure images as compared to a single long exposure.

Implementation of the ideal observer is further in the future, but its realization should be facilitated by the discussion in this paper.

Acknowledgments

This research was supported by Science Foundation Ireland under grant no. 01/PI.2/B039C and by an SFI Walton Fellowship (03/W3/M420) for H. H. Barrett. Development of the basic methodology for objective assessment of image quality was also supported in part by the National Institutes of Health under grant nos. R37 EB000803 and P41 EB002035.

References

1. Barrett, HH.; Myers, KJ. Foundations of Image Science. John Wiley and Sons; New York: 2004.
2. Barrett HH, Myers KJ, Devaney N, Dainty JC. Objective assessment of image quality: IV. Application to adaptive optics. 2006 Submitted to J. Opt. Soc. Am. A.
3. Hobson MP, McLachlan C. A Bayesian approach to discrete object detection in astronomical data sets. Mon Not R Astron Soc. 2003; 338:765784.
4. Braems, I.; Kasdin, NJ. Bayesian hypothesis testing for planet finding. 203rd Meeting; American Astronomical Society; 2004.
5. Barrett HH. Objective assessment of image quality: effects of quantum noise and object variability. J Opt Soc Am A. 1990; 7:1266–1278. [PubMed: 2370589]
6. Smith WE, Barrett HH. Hotelling trace criterion as a figure of merit for the optimization of imaging systems. J Opt Soc Am A. 1986; 3:717–725.
7. Fiete RD, Barrett HH, Smith WE, Myers KJ. The Hotelling trace criterion and its correlation with human observer performance. J Opt Soc Am A. 1987; 4:945–953. [PubMed: 3598746]
8. Hotelling H. The generalization of Student's ratio. Ann Math Stat. 1931; 2:360–378.
9. Khurd P, Gindi G. Decision strategies that maximize the area under the LROC curve. IEEE Trans Med Im. 2005; 24:1626–1636.
10. Veran JP, Rigaut F, Maitre H, Rouan D. Estimation of the adaptive optics long-exposure point-spread function using control loop data. J Opt Soc Am A. 1997; 14:3057–3068.
11. Aime C, Soummer R. The usefulness and limits of coronagraphy in the presence of pinned speckles. Astrophysical J. 2004; 612:L85L688.
12. Bloemhof EE. Anomalous intensity of pinned speckles at high adaptive correction. Opt Lett. 2004; 29:159–161. [PubMed: 14743997]
13. Perrin MD, Sivaramakrishnan A, Makidon RB, Graham JR. The structure of high strehl ratio point-spread functions. Astrophysical J. 2003; 596:702–712.
14. Sivaramakrishnan A, Lloyd JP, Hodge PE, Macintosh BA. Speckle decorrelations and dynamic range in speckle noise-limited imaging. Astrophysical J. 2002; 581:L59–L62.