

**Here is a sample chapter  
from this book.**

**This sample chapter is copyrighted  
and made available for personal use  
only. No part of this chapter may be  
reproduced or distributed in any  
form or by any means without the  
prior written permission of Medical  
Physics Publishing.**

# Wagner's Unified Theory of Image Quality: Three Decades Later

*Kyle J. Myers, Ph.D.,\* Iacovos S. Kyprianou, Ph.D., and Subok Park, Ph.D.  
Division of Imaging and Applied Mathematics  
Office of Science and Engineering Laboratories  
Center for Devices and Radiological Health (CDRH)  
U.S. Food and Drug Administration*

---

9.1	Introduction: R. F. Wagner and Image Quality	144
9.2	Elements of Wagner's Unified SNR for Medical Imaging	144
9.2.1	Classification Tasks and the Ideal Observer	144
9.2.2	The SNR of the Ideal Observer	145
9.3	Imaging as the Transfer of Information	145
9.3.1	Linear Systems Theory	146
9.3.2	Fourier Analysis for Linear Systems	147
9.3.3	Adding in Measurement Noise	147
9.3.4	Noise Equivalent Quanta (NEQ)	148
9.3.5	Strategy of the Ideal Observer	149
9.4	Extension to Shift-Variant, Non-Stationary Imaging	149
9.4.1	Object Representations	149
9.4.2	Singular Value Decomposition (SVD)	149
9.4.3	Noise Analysis for Digital Imaging Systems	152
9.4.4	SNR Transfer for Digital Systems	155
9.5	Current Trends	155
9.6	Conclusion	157
9.7	Acknowledgment	157
9.8	References	157

\*Corresponding author.

## 9.1 Introduction: R. F. Wagner and Image Quality

Over the course of his long and productive career, Robert F. Wagner revolutionized the field of medical imaging system evaluation. He was a brilliant electrical engineer and physicist with a remarkable intuition who synthesized ideas from World War II era communications theory, TV picture quality measures of the 1950s, and statistical decision theory publications from the 1960s to formulate a science of image quality for medical imaging. This chapter begins with an overview of the key concepts of Wagner's theory for medical image quality. In subsequent sections, we will consider various extensions to Wagner's theory that enable the evaluation of imaging system performance for modern system designs and applications. Bob Wagner's ideas are alive in the efforts of the current generation of imaging scientists, who continue to apply his approach to imaging system evaluation as they strive to design and optimize new approaches for disease detection and management through imaging.

## 9.2 Elements of Wagner's Unified SNR for Medical Imaging

Wagner's unified approach to the evaluation of medical imaging systems is founded on two core concepts (Wagner and Weaver 1972). The first is that image quality is *task dependent*. The only way to describe the quality of an imaging system is in terms of how useful it is for performing a specific task. An imaging system might be quite outstanding for one task, perhaps determining a calcium score, but not particularly effective for another, say detecting a lung nodule. Embedded in the task-dependent aspect of image quality is the requirement that the user of the images, that is, the strategy, algorithm, or person who performs the task, must be defined. Some imaging systems are targeted for use by highly trained experts, some for technologists, and still others might involve computer algorithms. The user of the images is termed the *observer*; in the next section we will take up the concept of optimal observers for the evaluation of imaging systems.

The second core concept in Wagner's theory is that image quality is *statistical*. Because all images are degraded by noise of one form or another, no two images of a particular patient are exactly the same. In fact, if we were able to acquire a large set of images of a single patient (imagining a world without radiation risks or other downsides to doing this), none of them would be exactly alike. Such a set of images from a single patient is random, where the randomness results from imaging system noise contributions (electrical, photon, thermal, etc.) and other sources of variability such as system or patient motion.

Consider a digital imaging system, such that each image comprises  $M$  pixels. An image can be represented mathemati-

cally as an  $M$ -dimensional vector  $\mathbf{g}$  with elements  $\{g_1, g_2, \dots, g_M\}$  for each of the  $M$  pixels. Here we use what is known as *lexicographic ordering* to create the data vector from the digital image. For example, if the image is the output of a 2D detector of dimension  $M_x$  by  $M_y$ , the first  $M_x$  elements of  $\mathbf{g}$  conventionally contain the data from the first row of the detector, the next  $M_x$  elements contain the second row, etc., until all  $M_y$  rows of the detector output are concatenated into one long data vector. The label  $m$  can take on any pixel address from  $(0,0)$  to  $(M_x, M_y)$ . In this way the pixel label  $m$  runs from 1 to  $M = M_x \times M_y$ , the total number of pixels or detected values in an image.

Patients, the objects at the input to the imaging system, can also be represented mathematically as vectors, although we have to keep in mind that they have very different mathematical properties from digital images. Objects, or patients, are essentially continuous. Further on in this chapter we will consider the computer generation of simulated images from models of objects and imaging systems. When such simulations are performed, the object is often modeled as a discretized or voxelized version of reality. Nevertheless, real objects are continuous, and we will represent them here by an infinite-dimensional vector  $\mathbf{f}$ .

The set of all images that could be obtained from a single object  $\mathbf{f}$  is described by a probability density function that describes the probability of obtaining an image  $\mathbf{g}$  given that the actual object was  $\mathbf{f}$ , written  $\text{pr}(\mathbf{g}|\mathbf{f})$ . Image quality is statistical because the observer's task performance depends very much on this probability density function. Image quality is defined in terms of expected performance of a task by an observer on average, given all the randomness in the image formation and capture processes. It is not meaningful to evaluate an imaging system on the basis of a single image, or to compare two systems using one image from each. Meaningful comparisons must take into account the nature of  $\text{pr}(\mathbf{g}|\mathbf{f})$ .

### 9.2.1 Classification Tasks and the Ideal Observer

When comparing imaging systems we must thus define the task and the observer. Wagner's work focused on *two-class detection* and *classification tasks*, often called binary tasks. In these tasks, the goal of the observer is to determine which of two hypotheses is true regarding the patient—for example, the tumor is either present or absent—on the basis of the image. This task presumes that each patient can be classified unambiguously one way or the other; there's no fuzziness or possibility that a patient belongs partially in both classes (and there's no option of belonging to some third category, either). Either hypothesis 1, denoted,  $H_1$ , is true, or else hypothesis 2, denoted  $H_2$ , is true. In a binary classification task, the observer must render a decision in favor of either  $H_1$  or  $H_2$  for each image.

One of Wagner's main contributions to medical imaging was his advocacy for the evaluation of imaging systems in

terms of how well the resulting images could be used for a classification task, when the task was performed by the best observer imaginable. He referred to this best-of-all-observers as the *ideal observer*. The ideal observer, a construct from statistical decision theory, is an observer that is optimal in just about any way imaginable. It is the observer that makes the fewest decision errors, it minimizes decision risk (when risks are assigned to each decision outcome), and it has the maximum hit rate at any false alarm rate.

The performance of the ideal observer gives an upper bound to the usefulness of the images from a specified system. To achieve this performance level, the ideal observer uses a classifier derived from a ratio of the likelihoods of the data under each of the hypotheses. The ideal observer must therefore have complete information regarding  $\text{pr}(\mathbf{g}|H_1)$ , the conditional probability of obtaining the set of images,  $\mathbf{g}$ , if  $H_1$  is true, and likewise  $\text{pr}(\mathbf{g}|H_2)$ , which requires knowledge of the image formation process (how data are governed by the underlying hypotheses, both deterministically as well as by any sources of randomness).

Wagner argued that this performance bound describes the full information content available in the images from that imaging system. The actual users of the images may not achieve this bound, but the information is there. Does that make the ideal observer irrelevant? Absolutely not, Wagner argued! Designers of imaging systems should make every effort to engineer systems that maximize the information available in the data for the intended task. It is then up to the actual users of the images, perhaps with the aid of computer-aided diagnosis software or image processing algorithms, to achieve performance approaching that bound.

### 9.2.2 The SNR of the Ideal Observer

In 1985, Wagner and Brown published a landmark paper on the “Unified SNR analysis of medical imaging systems.” In this work they introduced the signal-to-noise ratio of the ideal observer, denoted  $\text{SNR}_I$ . Wagner and Brown presented expressions for  $\text{SNR}_I$  for all the major medical imaging modalities of the time, including radiography, computed tomography (CT), nuclear magnetic resonance (NMR), and even time-of-flight PET (positron emission tomography), along with this simple expression as a *figure of merit* for the ideal observer's performance:

$$\text{SNR}_I^2 = \int d\boldsymbol{\rho} |\Delta S(\boldsymbol{\rho})|^2 \text{NEQ}(\boldsymbol{\rho}). \quad (9.1)$$

$\Delta S(\boldsymbol{\rho})$  is the Fourier transform or spectrum of the signal to be detected in a detection task ( $\boldsymbol{\rho}$  refers to spatial frequency), and its appearance in equation (9.1) ensures that  $\text{SNR}_I$  is task dependent: If the task is detection of a lesion,  $\Delta S(\boldsymbol{\rho})$  is the Fourier spectrum of the lesion, so it depends on the lesion's size, shape, amplitude, and so on. In a classification task,

$\Delta S(\boldsymbol{\rho})$  is the Fourier transform of the difference in the two objects at the heart of the competing hypotheses.

$\text{NEQ}(\boldsymbol{\rho})$  is the *noise equivalent quanta*—the number of quanta or photons the image is *worth*—to an ideal observer who makes perfect use of each quantum. Every medical image is made with a certain number of photons or counts of some sort, depending on the imaging modality. These photons represent in some sense the cost of the image, in terms of exposure to the patient, or time, or both. The more efficient the system is at “transferring” quanta, the closer the NEQ is to the actual number of quanta used to make the image. Thus, for the same input quanta, a better system is one that transfers quanta more efficiently, so that when the resulting  $\text{NEQ}(\boldsymbol{\rho})$  is weighted by the square of the spectrum of the difference between the signal alternatives being distinguished by the ideal observer, the better system has higher  $\text{SNR}_I$ . In the words of Wagner and Brown, “The ideal observer  $\text{SNR}_I^2$  is just an integral over the  $\text{NEQ}(\boldsymbol{\rho})$  spectrum determined by the system detection hardware—weighted by the spectrum of a difference signal—corresponding to the observer's task.”

It was not long after the introduction of this  $\text{SNR}_I$  approach that new radiological imaging systems were routinely described by their  $\text{NEQ}(\boldsymbol{\rho})$  in conference presentations, showing how the latest engineering innovations surpassed older generations of the system. Plots of  $\text{NEQ}(\boldsymbol{\rho})$ , or its close cousin, the detective quantum efficiency or  $\text{DQE}(\boldsymbol{\rho})$ , were routinely submitted (and still are) to the FDA as part of the laboratory data and system description provided by the manufacturer. The neat factorization of the task and system contributions to the  $\text{SNR}_I$  in equation (9.1) implied that systems could be compared by simply comparing their NEQs. The problem with that approach is that, as we shall see, the simply factorization of the Fourier quantities in equation (9.1) only holds when some very strong assumptions regarding the properties of the imaging system and the task are true.

For a meaningful discussion of the details of Wagner's unified SNR, and of its limitations, we have to take a few steps back and develop a framework for analyzing imaging systems in general. Then we will consider the various factors in Wagner's SNR in light of the known properties of modern digital imaging systems.

### 9.3 Imaging as the Transfer of Information

An imaging system is a mapping or transfer of information from an object, which we have denoted  $\mathbf{f}$ , to the image or data set, which we are calling  $\mathbf{g}$ . In the absence of measurement noise, the imaging process can be represented quite generally as (Barrett and Myers 2004):

$$\mathbf{g} = \mathcal{H}\mathbf{f}, \quad (9.2)$$

where  $\mathcal{H}$  is the deterministic operator that takes as input the object  $\mathbf{f}$  and generates a noiseless data set  $\mathbf{g}$ . This general expression allows for the imaging process to be either linear or nonlinear. In the sections that follow we will limit our analysis to linear imaging systems.

### 9.3.1 Linear Systems Theory

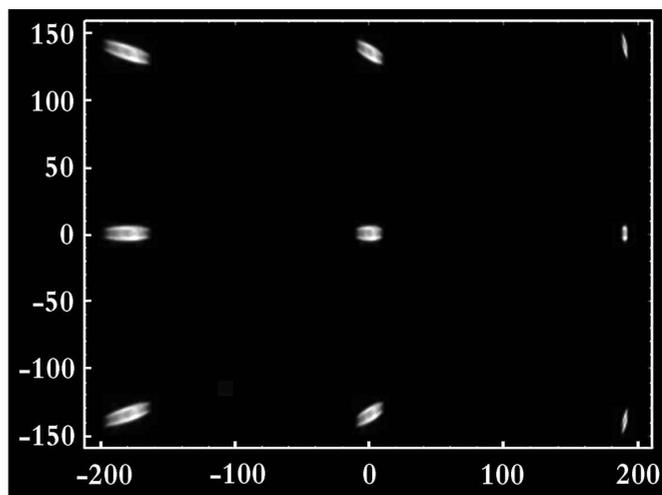
In Wagner's theory of image quality, the imaging system is assumed to be *linear* (or linearizable). In a linear system, if the object consists of two contributions,  $\mathbf{f} = \mathbf{f}_1 + \mathbf{f}_2$ , the expected image would be the sum of the images that would be obtained from each of the contributing objects:  $\mathbf{g} = \mathbf{g}_1 + \mathbf{g}_2 = \mathcal{H}\mathbf{f}_1 + \mathcal{H}\mathbf{f}_2$ . Likewise, if an object is composed of a signal and background, written  $\mathbf{f} = \mathbf{f}_s + \mathbf{f}_b$ , the image is the sum of the images of these contributions:  $\mathbf{g} = \mathcal{H}\mathbf{f}_s + \mathcal{H}\mathbf{f}_b = \mathbf{g}_s + \mathbf{g}_b$ . If a lesion's amplitude doubles, so would the resulting data from that signal. In many imaging modalities this assumption is quite close to being true, especially if the signal contribution is one of low contrast. Wagner was particularly interested in the detection performance of such systems in which signals are near the limit of detection, and in which contrast is thus inherently low.

When an imaging system can be treated as a linear operator, the mathematics describing the imaging process simplifies considerably. In that case the imaging system can be modeled in terms of *point response functions*,  $h(\mathbf{r}_{\text{det}}, \mathbf{r})$ . For a point source at location  $\mathbf{r}$  in the object,  $h(\mathbf{r}_{\text{det}}, \mathbf{r})$  describes the response of the imaging system at all locations  $\mathbf{r}_{\text{det}}$  in the detector plane. For an arbitrary object, system linearity tells us that the image is then a sum of the contributions from all the point sources that make up the object, as mapped by the point response functions (PRF):

$$g(\mathbf{r}_{\text{det}}) = \int h(\mathbf{r}_{\text{det}}, \mathbf{r}) f(\mathbf{r}) d\mathbf{r}. \quad (9.3)$$

The role of the PRF is suggested by the case of a perfect imaging system, in which  $h(\mathbf{r}_{\text{det}}, \mathbf{r})$  is represented by the Dirac delta function,  $h(\mathbf{r}_{\text{det}}, \mathbf{r}) = \delta(\mathbf{r}_{\text{det}} - \mathbf{r})$ . Then equation (9.3) reduces to  $g(\mathbf{r}) = f(\mathbf{r})$ , meaning that the transfer of information from the object to the image is exact.

Equation (9.3) assumes that while the system may not be perfect, it nonetheless is linear, and it allows for the PRF to depend on the location of the point in the object. This generality can be very important in some medical imaging modalities. In CT, for example, the response of the imaging system to a point in the center of the gantry is different from the image of one at the periphery. We call this system property *shift variance*. Similarly, in projection radiography the focal spot has a different effect on image blur depending on the location of the object of interest. As we see in Figure 9–1, the focal spot projection onto a detector through a pinhole positioned at different locations in object space has different



**Figure 9–1.** Projection of the focal spot from a radiographic imaging system onto different points on the detector plane, demonstrating the location-specific differences in blur than can result.

shape and size (Kyprianou et al. 2006). This effect causes the blur of structures near the chest wall in mammography to be different from that experienced by structures far from the chest wall, for example.

A full description of a shift-variant imaging system requires the characterization of  $h(\mathbf{r}_{\text{det}}, \mathbf{r})$  for all combinations of points in object space and locations in the resulting image. In some applications this can be done experimentally, as for example the approach pioneered at the Center for Gamma Ray Imaging at the University of Arizona. In that group's approach, a small radioactive point source, mounted on a robotic arm, is stepped through a large number of locations inside the bore of a single photon emission computed tomography (SPECT) system, allowing the determination of the response of the system to each "object point" location (Rowe et al. 1993). This information is stored for use in system-specific image reconstruction algorithms and estimation routines.

Another approach to the characterization of a shift-variant imaging system is through the use of accurate Monte Carlo modeling. With such simulations, the PRF of a modern digital imaging system can be probed in detail at very fine scale. In radiography, angle of incidence is another variable that needs to be considered for modern geometries, especially in cone-beam and tomosynthesis systems. Badano and coworkers at the FDA's Center for Devices and Radiological Health (CDRH) have published examples of the variation in the point response as a function of the angle of incidence in a digital radiography system (Badano et al. 2006; Kyprianou et al. 2008). An important challenge in making use of such detailed simulations is the development of parsimonious representations of the resulting data for system modeling, comparisons of competing system designs, and improved image reconstruction algorithms and inferences from the resulting images (Rao et al. 2010).

If we are able to assume that the PRF is the same for all object locations, we say that the system is *shift invariant*. The PRF for a *linear, shift-invariant (LSIV)* imaging system is referred to as the *point spread function (PSF)*. For an LSIV system, the imaging process is modeled as a simple convolution of the PSF with the object:

$$g(\mathbf{r}_{\text{det}}) = \int h(\mathbf{r}_{\text{det}} - \mathbf{r}) f(\mathbf{r}) d\mathbf{r}. \quad (9.4)$$

In other words, there is no preferred origin of the imaging system. The image of any point source in the object is the same—it has the same point spread function—with the requisite shift to account for the point source's location. The image of a complicated object is then realized as a sum of shifted and weighted PSFs.

An LSIV system must have a continuous data set. Digital imaging systems are never truly LSIV. Because of the finite size of each detector element in a digital imaging system, slight shifts in an object location can result in large changes in the data. Thus the imaging equation in (9.4) may be a good model for a film-screen x-ray system, but not nearly so for an x-ray system with a digital detector, especially one with large pixels.

### 9.3.2 Fourier Analysis for Linear Systems

From the very beginning of his career, Bob Wagner championed the use of Fourier analysis as a tool for characterizing medical imaging systems. In 1972, not long after he arrived at the FDA, he published a review of image quality metrics in an SPIE proceedings document (see also Brown 2009). In this early analysis, Wagner laid out his vision for task-based image quality assessment and the Fourier quantities central to the approach he advocated. It is these quantities that we take up next.

The assumption that an imaging system acts as a linear mapping is a key enabler of using Fourier methods to analyze its properties. With Fourier analysis, a continuous object is written as a weighted sum of Fourier basis functions. In two dimensions, the continuous object  $f(\mathbf{r})$ , or  $f(x, y)$ , can be expressed as:

$$f(x, y) = \int_{-\infty}^{\infty} d\xi \int_{-\infty}^{\infty} d\eta F(\xi, \eta) \exp(2\pi i \xi x + 2\pi i \eta y). \quad (9.5)$$

The Fourier basis functions are the complex exponentials,  $\exp(2\pi i \xi x + 2\pi i \eta y)$ , which are simply two-dimensional wave functions with spatial frequencies  $\xi$  and  $\eta$  in the  $x$  and  $y$  directions, respectively. Each basis function has weight  $F(\xi, \eta)$ , which we refer to as the Fourier coefficients of the object.

Equation (9.5) is the familiar form of the inverse Fourier transform for the specific case of a two-dimensional object.

Since in medical imaging the object might be three-dimensional, as in the case of volumetric imaging, or even four-dimensional when time-dependent properties of the object are important, we can express equation (9.5) more generally as

$$f(\mathbf{r}) = \int F(\boldsymbol{\rho}) \exp(2\pi i \mathbf{r} \cdot \boldsymbol{\rho}) d\boldsymbol{\rho}, \quad (9.6)$$

where  $\mathbf{r}$  is the vector of coordinates in the object space and  $\boldsymbol{\rho}$  is the corresponding conjugate vector of frequencies in the Fourier domain. The integral is over the infinite range of frequencies.

Like all good basis functions, the Fourier basis has the beautiful property of being *orthonormal*:

$$\int \exp(2\pi i \mathbf{r} \cdot \boldsymbol{\rho}_1) \exp(-2\pi i \mathbf{r} \cdot \boldsymbol{\rho}_2) d\mathbf{r} = \delta(\boldsymbol{\rho}_1 - \boldsymbol{\rho}_2), \quad (9.7)$$

where the right-hand side features the Dirac delta function. The integrated product of any one basis function with itself gives 1, while the integrated product of any basis function with a different basis function gives 0. Here again the integrals cover the infinite range, this time over all space.

The orthonormality property of the Fourier basis set is what allows us to model an LSIV imaging system as a transfer of independent frequencies through the system. Fourier transformation of the LSIV imaging equation (9.4) gives:

$$G(\boldsymbol{\rho}) = H(\boldsymbol{\rho}) F(\boldsymbol{\rho}), \quad (9.8)$$

where  $H(\boldsymbol{\rho})$ , the Fourier transform of the system PSF, is referred to as the *transfer function* of the system.  $H(\boldsymbol{\rho})$  is related to the *modulation transfer function (MTF)* of the imaging system according to  $\text{MTF}(\boldsymbol{\rho}) = |H(\boldsymbol{\rho})|/H(0)$ . As its name suggests, the MTF of an imaging system is a description of the ability of a system to transfer modulations, or image contrast if you wish, as a function of spatial frequency from input to output. By its normalization, the MTF ranges from 0 to 1;  $\text{MTF}(\boldsymbol{\rho}) = 1$  indicates that an image's Fourier component at spatial frequency  $\boldsymbol{\rho}$  is passed through the system without decrease in amplitude. Alternatively, an  $\text{MTF}(\boldsymbol{\rho})$  of 0 means the information being conveyed at that frequency is completely attenuated by the system, so that nothing is measurable at the output. In Fourier-speak, equation (9.8) represents shift-invariant filtering of the object.

### 9.3.3 Adding in Measurement Noise

All imaging systems suffer from one or more sources of measurement noise. We can express this reality by adding a noise term to the imaging equation:

$$\mathbf{g} = \mathcal{H}\mathbf{f} + \mathbf{n} \quad (9.9)$$

Expressing the noisy image in this way does not restrict our analysis to systems in which the noise is additive. Systems limited by quantum noise, which is known to be Poisson and thus image-dependent (certainly not additive), can be written in this way if we understand equation (9.9) to mean that the noise is the difference between the actual image and the image we expected in the absence of noise:  $\mathbf{n} = \mathbf{g} - \mathcal{H}\mathbf{f}$ .

Because of measurement noise, an image is a sample of a random process. A full description of the data's randomness would require knowledge of the full probability density of the data given the object,  $\text{pr}(\mathbf{g}|\mathbf{f})$ . This probability density is often modeled using familiar forms from probability theory. For example, we often assume the data are Gaussian, as in the case of thermal noise fluctuations from electronic circuits and MR receiver coils, or we may adopt a Poisson model, as in quantum-limited imaging. These probability laws describe the fluctuations in the data values in each pixel, the correlations between pixels, and all higher-order statistical relationships between the random data values.

We use simpler descriptions of the noise, especially low-order moments of the data such as their mean value and standard deviation, for several reasons. First, if it is reasonable to assume that the noise follows some well-known distribution like the Poisson or Gaussian model, then such low-order moments completely describe it. A Poisson distribution is completely characterized by the mean of the random parameter, and a Gaussian data vector is completely described by the mean data vector and its covariance. The other reason we use low-order moments to describe random images is that we are able to estimate them from a manageable number of measured images. We will have more to say on this practical issue later in this chapter. For now, we assume we have an infinite population of images, or equivalently, complete information regarding the randomness in the images. Furthermore, we'll assume for now that the noise is additive and Gaussian distributed. Wagner's original contributions on image quality presumed this form of noisy data, arguing that multiplicative Poisson noise is adequately modeled as Gaussian in the limit of low-contrast imaging. At the same time his analysis accounted for correlations in the images.

The use of Fourier analysis to describe image noise correlations is predicated on certain powerful assumptions regarding the behavior of the underlying random process. In a few words, the random process is assumed to be *wide-sense stationary*. Stationarity in a broad sense stipulates that the randomness of the data does not depend on origin or absolute location. Wide-sense stationarity is a weak form of stationarity: the randomness of the data is only required to be independent of absolute location when considering the mean and the covariance of the data.

Under wide-sense stationarity, the average value of a set of images is the same everywhere. If  $\langle g(\mathbf{r}) \rangle$  is the average of an infinite number of images, wide-sense stationarity requires

that  $\langle g(\mathbf{r}) \rangle = C$ ; the average image is a uniform field of level  $C$ .

In the analysis of image noise, it is useful to separate out the difference in the mean at each image location from the fluctuations about that mean, in which case the *autocovariance function* becomes the statistical description of choice. The autocovariance function,  $K(\mathbf{r}_1, \mathbf{r}_2)$ , is defined by:

$$K(\mathbf{r}_1, \mathbf{r}_2) = \left\langle \left[ g(\mathbf{r}_1) - \bar{g}(\mathbf{r}_1) \right] \left[ g^*(\mathbf{r}_2) - \bar{g}^*(\mathbf{r}_2) \right] \right\rangle, \quad (9.10)$$

where the overbars denote average quantities and the asterisk denotes complex conjugate. The autocovariance function describes the covariation in the images between locations  $\mathbf{r}_1$  and  $\mathbf{r}_2$ . A full characterization of the autocovariance would require knowledge of this function for all pairs of locations in the image. Wide-sense stationary processes are ones in which the autocovariance does not depend on absolute location, but rather depends only on the vector separation  $\Delta\mathbf{r} = \mathbf{r}_2 - \mathbf{r}_1$ . In that case

$$K(\Delta\mathbf{r}) = \left\langle \left[ g(\mathbf{r}) - \bar{g}(\mathbf{r}) \right] \left[ g^*(\mathbf{r} + \Delta\mathbf{r}) - \bar{g}^*(\mathbf{r} + \Delta\mathbf{r}) \right] \right\rangle. \quad (9.11)$$

While the full autocovariance function of equation (9.10) depends on two position vectors, the wide-sense form of equation (9.11) only depends on one, the vector separation  $\Delta\mathbf{r}$  between the two points. The absolute location of the two points is immaterial (there is no preferred origin). This is the *stochastic* data description's analogy to shift-invariance, which we described above as meaning that the *deterministic* image formation process can be written in terms of a point response that depends only on the distance from the center of the point's image, not on the absolute location of that image.

The standard Fourier description of image noise is the *Wiener* or *noise power spectrum*,  $\text{NPS}(\boldsymbol{\rho})$ , a frequency-dependent description of the fluctuations in the output of the imaging system. The noise power spectrum is related to the autocovariance function via the Fourier transform  $\mathcal{F}$ :

$$\left[ \mathcal{F} K \mathcal{F}^\dagger \right](\boldsymbol{\rho}_1, \boldsymbol{\rho}_2) = \text{NPS}(\boldsymbol{\rho}_1) \delta(\boldsymbol{\rho}_1 - \boldsymbol{\rho}_2), \quad (9.12)$$

where the dagger implies an adjoint (see more on adjoints below). This expression relies on the wide-sense stationarity of the image noise. It is only in that case that the autocovariance is a function of a single vector, which can be transformed to yield a noise power spectrum that is also a function of one (spatial frequency) vector. In later sections we will consider noise descriptions for situations where the data are not stationary.

### 9.3.4 Noise Equivalent Quanta (NEQ)

We now have all the Fourier building blocks needed to define the Noise Equivalent Quanta of equation (9.1).  $\text{NEQ}(\boldsymbol{\rho})$  is

the product of three Fourier descriptions of an imaging system's performance: the grayscale transfer  $H(0)$ , the MTF, and the NPS, according to

$$\text{NEQ}(\boldsymbol{\rho}) = \frac{H^2(0) \times \text{MTF}^2(\boldsymbol{\rho})}{\text{NPS}(\boldsymbol{\rho})}. \quad (9.13)$$

When we substitute this expression into the  $\text{SNR}_I$  of equation (9.1), we see that the ideal observer's figure of merit has large contributions from spatial frequencies in the signal that are transferred efficiently by the system (large values of the MTF), while those frequencies that are attenuated by the system (low MTF) contribute less. Spatial frequencies for which the NPS is high bring down the  $\text{SNR}_I$ , relative to those frequencies for which the noise is low.

### 9.3.5 Strategy of the Ideal Observer

The  $\text{SNR}_I$  of equation (9.1) is achievable in practice by an algorithm or machine reader known as the *prewhitening matched filter*. This ideal observer applies a linear filtering algorithm to the Fourier transform of each image,  $G(\boldsymbol{\rho})$ , to give the following classifier output:

$$t = \int d\boldsymbol{\rho} \frac{\Delta \bar{G}(\boldsymbol{\rho}) G(\boldsymbol{\rho})}{\text{NPS}(\boldsymbol{\rho})} = \int d\boldsymbol{\rho} \frac{H(\boldsymbol{\rho}) \Delta S(\boldsymbol{\rho}) G(\boldsymbol{\rho})}{\text{NPS}(\boldsymbol{\rho})}, \quad (9.14)$$

where  $\Delta \bar{G}(\boldsymbol{\rho}) = H(\boldsymbol{\rho}) \Delta S(\boldsymbol{\rho})$  is the Fourier transform of the image of the expected difference in the objects; if the task is signal detection, then this is the image of the lesion to be detected. A computer-aided diagnosis algorithm is sensitive to image features; likewise, the ideal observer makes use of features, in this case, an infinite sum of them, evaluated at spatial frequencies in the Fourier domain. Each feature in frequency space is the product of the Fourier transform of the image with the Fourier transform of the image of the known signal to be detected, divided by the noise power in that frequency. Early on, owing to his training as an electrical engineer, Bob Wagner thought of those frequencies as channels.

## 9.4 Extension to Shift-Variant, Non-Stationary Imaging

In modern medical imaging, the object is continuous while the data are digital, which is to say, discrete. Thus, the imaging system is mathematically a continuous-to-discrete (CD) operator. Moreover, the digital detector has a finite extent. So the assumptions that led to the Fourier form of the ideal observer's  $\text{SNR}_I$  in equation (9.1) are challenged by modern medical imaging systems.

Nevertheless, as we shall see, the guiding principles that led to that figure of merit are very much applicable to the analysis of today's imaging modalities. In particular, image

quality must be objective and task-based, requiring that the object(s) to be detected/classified/estimated are specified, as must also be the observer that will perform the task. Second, the objective assessment of image quality requires an accurate accounting of the deterministic and stochastic properties of the image formation process. And finally, as we shall see, the tools available to us from linear systems theory are powerful resources for analyzing even systems with properties that call into question standard Fourier analysis.

### 9.4.1 Object Representations

In section 9.3.2 we wrote the continuous object in terms of a weighted sum of wavefunctions, equation (9.5). The Fourier basis is just one choice for basis functions available for representing objects. More generally, we can *exactly* represent *any* object as an infinite sum of weighted orthonormal basis functions  $u_n(\mathbf{r})$ :

$$f(\mathbf{r}) = \sum_{n=1}^{\infty} \alpha_n u_n(\mathbf{r}) \quad (9.15)$$

as long as the object has finite extent and "energy" (the integral of the object squared). The reader may note a difference between this expansion representation and the integral form in equation (9.5). The difference comes from the acknowledgment here that real patients have finite size. The Fourier expression of equation (9.5) allowed for objects that had no spatial boundary.

The orthonormality property of basis functions (refer to equation (9.7)), makes the determination of the coefficients in equation (9.15) straightforward:

$$\alpha_n = \int u_n^*(\mathbf{r}) f(\mathbf{r}) d\mathbf{r}. \quad (9.16)$$

So now the question arises, how do we choose the basis set  $u_n(\mathbf{r})$ ? The answer is inspired by this fact: for LSIV systems, the Fourier wavefunctions have the special property that they are *eigenfunctions* of the imaging system operator. That is, the system can be seen as transferring the Fourier wavefunctions independently (the transfer of one frequency of the object is not influenced by the presence of another) according to the system transfer function for that frequency. By extension, we seek a basis that provides the same ability to model the digital image acquisition process as the transfer of independent basis functions.

### 9.4.2 Singular Value Decomposition (SVD)

In the absence of noise, a linear continuous-to-discrete imaging system generates data according to:

$$g_m = \int h_m(\mathbf{r}) f(\mathbf{r}) d\mathbf{r}, \quad (9.17)$$

where  $g_m$  is the data in the  $m^{\text{th}}$  pixel and  $h_m(\mathbf{r})$  is the  $m^{\text{th}}$  point response function, which describes the sensitivity of the  $m^{\text{th}}$  measurement to the function  $f(\mathbf{r})$  at point  $\mathbf{r}$ . Here we use the process of lexicographic ordering described in section 9.2 to create a vector of  $M$  elements from the digital image. Equation (9.17) describes a CD mapping that results in a tremendous dimensionality reduction. At the input is an object that requires an infinite set of coefficients to represent exactly (cf. equation (9.15)); the output is a finite set of measurements,  $M$  of them in all. Clearly this is not an invertible process. There is simply no way to determine the true object that was at the input of the imaging system from its digital image  $\mathbf{g}$ . Mathematically,  $\mathbf{g} = \mathcal{H}\mathbf{f}$ , but we cannot write  $\mathbf{f} = \mathcal{H}^{-1}\mathbf{g}$  because  $\mathcal{H}^{-1}$  does not exist.

This noninvertibility of the imaging operator, owing to the continuous-to-discrete nature of digital imaging systems, is a fact we did not reckon with when we were considering LSIV systems. There we assumed that the dimensionality of the object and image were the same. Both were continuous; the imaging system was a continuous-to-continuous (CC) mapping. So long as the system MTF had no zeroes, the object could be recovered by inverting equation (9.8). While any real CC system will have some frequencies for which the MTF is zero, it is at least possible to see from equation (9.8) how to recover the transferred frequencies of the object.

As a stepping stone to determining the basis functions we seek, we introduce an alternative operator that maps the discrete data to a function in the space of the object. The adjoint operator  $\mathcal{H}^\dagger$  takes as its input a vector in the data space and outputs  $\tilde{f}(\mathbf{r})$ , a superposition of the complex conjugates of the point response functions  $h_m(\mathbf{r})$ :

$$\tilde{f}(\mathbf{r}) = [\mathcal{H}^\dagger \mathbf{g}](\mathbf{r}) = \sum_{m=1}^M g_m h_m^*(\mathbf{r}). \quad (9.18)$$

The function  $\tilde{f}(\mathbf{r})$  lives in the space of the object. Operationally  $\tilde{f}(\mathbf{r})$  is a weighted sum of the continuous functions  $h_m^*(\mathbf{r})$ , where the weights are the data elements. The adjoint operator exists and is well behaved, even when the same cannot be said of the inverse operator.

The adjoint operator converts a function of image space coordinates to a function of object space coordinates. Its behavior is familiar for special cases. For example, the adjoint of a real matrix is its transpose. In projection imaging, the adjoint of the continuous-to-discrete operator describing the image formation process is the *backprojection* operator. The adjoint of the Fourier transform operator is the inverse Fourier transform. (The adjoint of any unitary operator is its inverse; the Fourier operator is one example of this general property.)

It is tempting to drop the complex conjugate in equation (9.18), since for many applications, including x-ray and

gamma-ray imaging, the objects are real, the data are real, and the point response functions are real. However, there are other medical imaging applications where the data are complex, most notably MRI and ultrasound, such that it is important to retain the complex conjugate notation.

Having defined the adjoint operator, we can now form two new operators,  $\mathcal{H}^\dagger \mathcal{H}$  and  $\mathcal{H} \mathcal{H}^\dagger$ . The former maps a function in object space to another function in object space, for example, via projection followed by backprojection. The  $\mathcal{H} \mathcal{H}^\dagger$  operator maps a vector in the space of the digital data to another digital vector of the same dimensionality. Thus both these new operators have input and output functions that belong to the same space.  $\mathcal{H}^\dagger \mathcal{H}$  is a continuous-to-continuous operator;  $\mathcal{H} \mathcal{H}^\dagger$  is a discrete-to-discrete operator (a matrix).

Recall that for continuous-to-continuous LSIV imaging systems we were able to analyze the system in terms of wavefunctions that are transferred independently by the system. Our goal here in making these operator “sandwiches” is to construct new operators with the right array of properties such that they are amenable to a similar analysis in terms of eigenvectors and eigenvalues of the system.

The eigenvector-eigenvalue equation for  $\mathcal{H}^\dagger \mathcal{H}$  is given by

$$\mathcal{H}^\dagger \mathcal{H} \mathbf{u}_n = \mu_n \mathbf{u}_n, \quad (9.19)$$

where the  $\mathbf{u}_n$  are the eigenvectors of the compound operator and their associated eigenvalues are the  $\mu_n$ . Equation (9.19) tells us that the basis functions  $\mathbf{u}_n$  are transferred by the operator  $\mathcal{H}^\dagger \mathcal{H}$  via a change in amplitude only; they are unchanged in form.

All of the eigenvalues  $\mu_n$  are non-negative, although some may be zero. The number of nonzero eigenvalues is the *rank*,  $R$ , of the eigensystem. The eigenvectors associated with a nonzero eigenvalue are said to be in the *measurement space* of the system. Those eigenvectors that have an eigenvalue of zero are said to be *null functions* of the imaging system. Any arbitrary object can be written as an expansion (cf. equation (9.15)) using the basis functions that satisfy the eigensystem equation (9.19). We can write equation (9.15) as two sums, one over the expansion subset associated with nonzero eigenvalue, and the other over the expansion associated with an eigenvalue equal to zero:

$$\mathbf{f} = \mathbf{f}_{meas} + \mathbf{f}_{null} = \sum_{n=1}^R \alpha_n \mathbf{u}_n + \sum_{n=R+1}^{\infty} \alpha_n \mathbf{u}_n, \quad (9.20)$$

where the subscript *meas* refers to the measurement space of the imaging operator, and the subscript *null* refers to the null space of the imaging system. The null space is the space of all objects that, when imaged, yield no data:  $\mathcal{H}\mathbf{f}_{null} = 0$ . Another way of saying this is that an arbitrary number of null functions can be added to an object and the data will be unchanged *in the absence of noise*.

Every digital imaging system has null functions. It is the existence of null functions that causes a digital imaging system to be noninvertible. Because a null function is not measurable, the data can never provide sufficient information to determine which actual object was at the input to the system. They are the counterpart of the spatial frequencies of an LSIV system for which the MTF is zero.

Imaging scientists are quite familiar with the ramifications of eigenanalysis for LSIV systems. An LSIV system transfers a complex exponential wavefunction of frequency  $\rho$  unchanged except for scaling by the eigenvalue  $H(\rho)$ , as described by equation (9.8). What is less familiar to imaging scientists is the notion that *any* linear system can be subject to eigenanalysis by the simple expediency of creating an eigensystem that can be analyzed by sandwiching together  $\mathcal{H}$  and  $\mathcal{H}^\dagger$  as we have done here.

We can write a related expression to equation (9.19) for the eigenvectors of  $\mathcal{H}\mathcal{H}^\dagger$ :

$$\mathcal{H}\mathcal{H}^\dagger \mathbf{v}_n = \mu_n \mathbf{v}_n \quad (9.21)$$

where the  $\mathbf{v}_n$  form an orthonormal set of basis functions for the digital data. The notation in equations (9.19) and (9.21) is intended to make it clear that the same eigenvalues are found when analyzing the eigenequations for both  $\mathcal{H}\mathcal{H}^\dagger$  and  $\mathcal{H}^\dagger\mathcal{H}$ .

The basis vectors  $\{\mathbf{u}, \mathbf{v}\}$  and coefficients  $\mu_n$  are known as the *singular system* for the imaging operator  $\mathcal{H}$ . The general form of that linear operator can be expressed as:

$$\mathcal{H} = \sum_{n=1}^R \sqrt{\mu_n} \mathbf{v}_n \mathbf{u}_n^\dagger \quad (9.22)$$

*Singular vector decomposition (SVD)* is the formal name for the analysis we have just laid out. With the SVD decomposition of the imaging operator in equation (9.22) and that of the object in equations (9.15) and (9.20), the imaging process can be rewritten as:

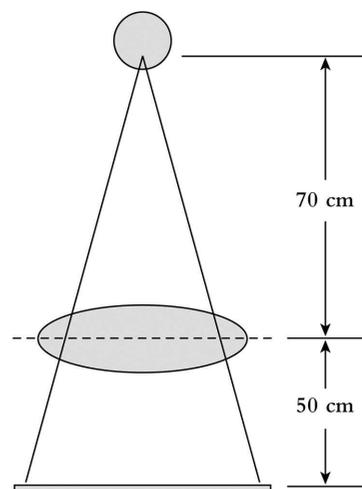
$$\mathbf{g} = \mathcal{H}\mathbf{f} = \sum_{n=1}^R \sqrt{\mu_n} \mathbf{v}_n \mathbf{u}_n^\dagger \sum_{m=1}^{\infty} \alpha_m \mathbf{u}_m = \sum_{n=1}^R \alpha_n \sqrt{\mu_n} \mathbf{v}_n = \sum_{n=1}^R \beta_n \mathbf{v}_n \quad (9.23)$$

The double sum collapses to a single sum because of the orthonormality of the  $\mathbf{u}_n$ . All that remains is a sum over the  $\mathbf{v}_n$  with coefficients  $\beta_n = \alpha_n \sqrt{\mu_n}$ . Thus the data set for a particular object can be easily found: determine the coefficients for that object using SVD, and multiply them by the square root of the transfer values for those basis vectors to determine the weights for the basis vectors in data space. SVD turns an arbitrary linear mapping into a simple multiplication. We

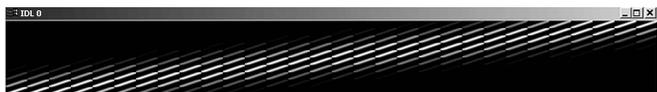
have gotten what we were looking for in terms of an information transfer way of thinking of our digital imaging system.

Standard numerical recipes exist for matrix multiplication and eigenanalysis of linear systems. Figure 9–2 illustrates the concept with a toy problem. Here we simulate a radiographic imaging system with a Gaussian focal spot, a digital detector with 0.2 mm pixels, Gaussian point spread function, and pixel gain nonuniformity of  $\pm 2.5\%$ . The system geometry is set up to give a magnification of 1.7. No scatter is included in the model. We use simple ray-tracing to calculate the point response functions for a large number of finely sampled points in the object such that the images of these points will be centered at the vertices of the rectilinear grid at the detector, at the centers, and at points in between. In this way we are able to approximate a continuous-to-discrete imaging operator.

For each point object, we calculate an image according to equation (9.17). We then turn that two-dimensional image into a one-dimensional column vector using the lexicographic ordering process we described in section 9.2. Figure 9–3 is a matrix containing the images of all the point locations we sampled, where each column is the image of a point at a different location. The process can easily be reversed, putting each column back into a square format to visualize it as an image. We can see that the columns (the images of the various points) are all quite similar, although there is shifting of the bright values that happens as the point sources shift location in object space. The images also differ slightly because the pixels in the detector have non-uniform gain. Because we computed PRFs for a large number of finely spaced object locations (to sample subpixel shifts in the detector space), the matrix is very nonsquare. The number of columns equals the number of point sources we “imaged.” The number of rows



**Figure 9–2.** Geometry of imaging system for example singular system analysis. SEE COLOR PLATE 15.



**Figure 9-3.** Matrix containing PRFs for system shown in Figure 9-2. SEE COLOR PLATE 16.

equals the size of the image format. This matrix is a simulated representation of the system operator  $\mathcal{H}$ .

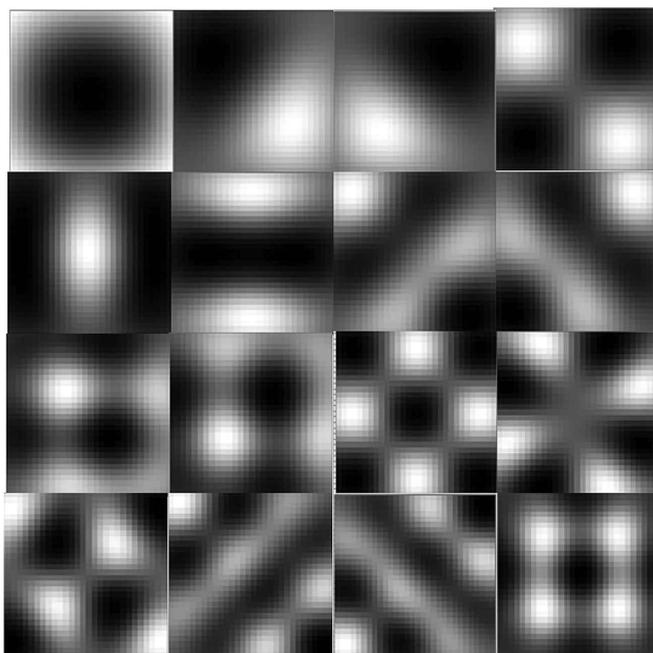
Figure 9-4 is the compound operator  $\mathcal{H}^\dagger \mathcal{H}$ , obtained by multiplying the matrix in Figure 9-3 by its transpose. We see that this matrix is square. This is the matrix we feed into a numerical algorithm to determine its eigenvectors and eigenvalues. These, in turn, fully describe the singular system for  $\mathcal{H}$ .

Figure 9-5 shows the first 16 singular vectors  $\mathbf{u}_n$  of the imaging system, depicted as quasi-continuous (simulated) objects. Each of these vectors in object space is transferred by the imaging system according to equation (9.23), that is, by a simple amplitude scaling, where the scale factor is the associated  $\sqrt{\mu_n}$ .

Any arbitrary object can be decomposed into a weighted sum of the discovered singular vectors according to equation (9.15). To determine the weights for each vector, we simply compute the inner product of the arbitrary object with each singular vector, according to equation (9.16). The object will be found to be the sum of its measurement (the singular vectors with nonzero  $\mu_n$ ) and null components (the singular vectors with  $\mu_n = 0$ ) according to equation (9.20).



**Figure 9-4.** Compound operator  $\mathcal{H}^\dagger \mathcal{H}$ . SEE COLOR PLATE 17.



**Figure 9-5.** First 16 eigenvectors of the matrix shown in Figure 9-4. SEE COLOR PLATE 18.

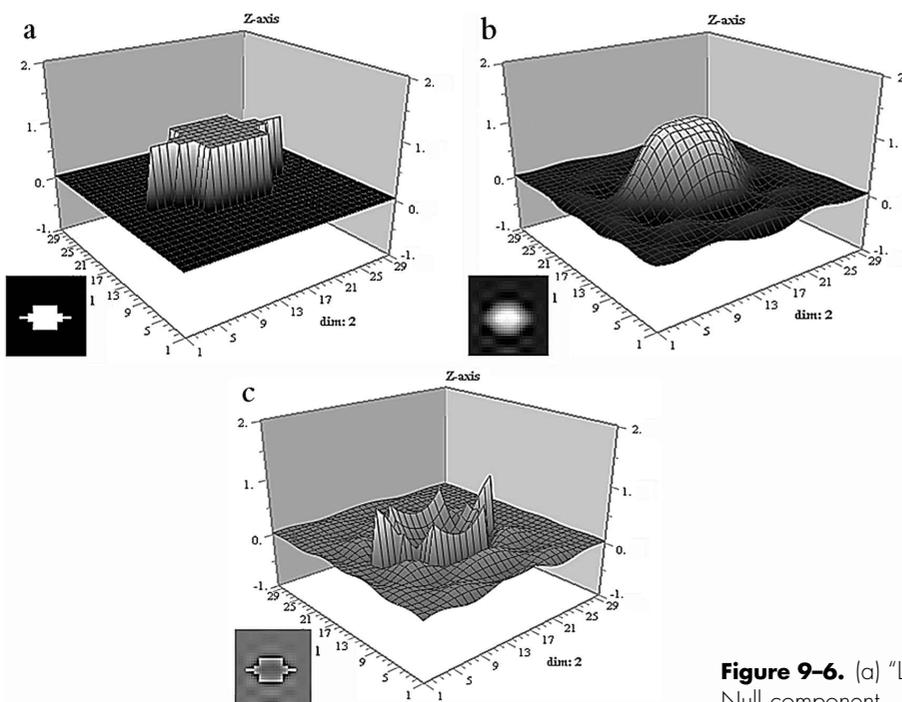
As an illustration, consider the simulated “lesion” shown in Figure 9-6a. When we determine its measurement and null components, we find that the object is the sum of the vectors shown in Figures 9-6b and 9-6c. The noise-free image of the simulated lesion will be equal to the noise-free image of the measurement component shown in Figure 9-6b. Note that the measurement component lacks much of the detail of the true lesion. This detail cannot contribute to the data, even in the absence of noise.

By determining the measurement space of an imaging system, we can get an understanding of the limits of the system’s performance in the absence of noise. In the toy problem illustrated in Figures 9-2 through 9-6, we focused our attention on a thin or planar object. The system model did not consider attenuation or scatter, as would occur in the presence of a thick object. Our simplified illustration provides in some sense an upper bound on the expected performance of the imaging system for imaging a planar object embedded in a thick medium. See also Kyprianou et al. 2008 and 2009 and Liu et al. 2009 for further examples of eigensystems of digital imaging systems.

#### 9.4.3 Noise Analysis for Digital Imaging Systems

For continuous data, the noise can be described by the auto-covariance function as given in equation (9.10). For discrete data, the counterpart is the *covariance matrix*:

$$\mathbf{K} = \langle (\mathbf{g} - \bar{\mathbf{g}})(\mathbf{g} - \bar{\mathbf{g}})^\dagger \rangle \quad (9.24)$$



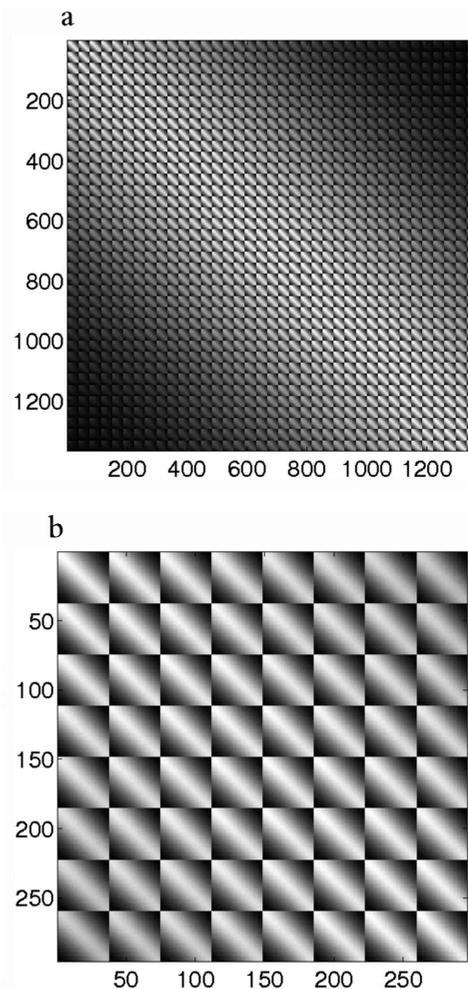
**Figure 9-6.** (a) “Lesion” object. (b) Measurement component. (c) Null component.

with matrix elements

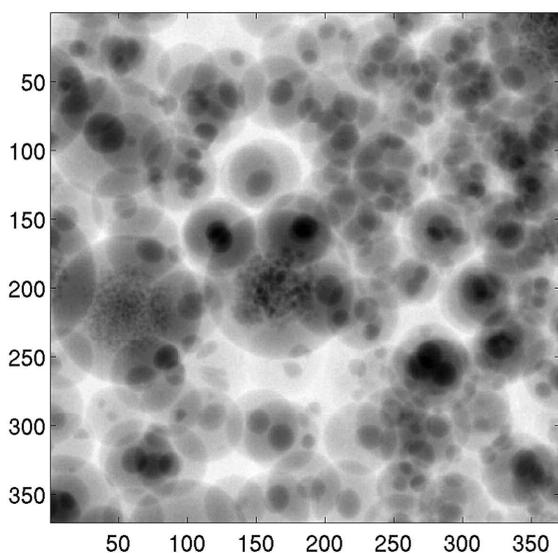
$$K_{n,m} = \langle (g_n - \bar{g}_n)(g_m - \bar{g}_m)^* \rangle. \quad (9.25)$$

The covariance matrix is square and of dimension  $M \times M$ . The pixel variances  $\sigma_m^2 = \langle (g_m - \bar{g}_m)^2 \rangle = K_{m,m}$  run down the diagonal of  $\mathbf{K}$ . The off-diagonal elements are the covariances between the different pixels in the image. This description is fully general, allowing for nonstationary random vectors. A stationary data set is one for which  $K_{n,m} = K_{n-m}$ .

Figure 9-7a shows the covariance matrix for the data from a set of ball phantom images like the one in Figure 9-8. The bright central diagonal is the variance of each pixel with itself. The next strongest contribution to the covariance matrix is the correlation of each pixel in the data with its nearest neighbor; this is seen in the bright diagonal lines on either side of the center diagonal. And so on. The covariance matrix is magnified in Figure 9-7b to show the tile structure found within it. This structure results from the lexicographic ordering used to assign the pixel elements in the detector to the long data vector  $\mathbf{g}$  used to generate the covariance matrix  $\mathbf{K}$ . Figure 9-8 is an x-ray image of a container filled with a random collection of plastic balls of various sizes and densities (Park et al. 2009a). If we take a collection of images like this, with the balls stirred up before each acquisition to randomize their locations, we have a set of random images  $\{\mathbf{g}\}$ . The randomness in the data will be the result of measurement noise and the random variation in the locations of the balls. In the same way, we can imagine having a set of patient images with varying patient structure and image noise.



**Figure 9-7.** (a) Covariance matrix of the random ball phantom. (b) Zoomed in version of the image in (a).



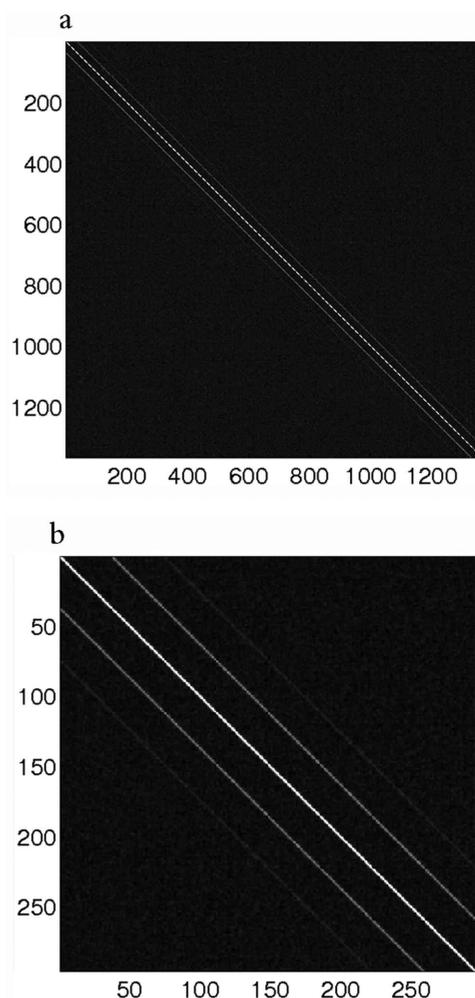
**Figure 9-8.** X-ray image of phantom containing random assortment of balls of range of sizes and densities.

The structure in the covariance in Figure 9-7 shows the long-range correlations in the data from the balls in the object. For comparison, Figure 9-9 shows the covariance matrix for a collection of images of a uniform phantom, in this case, a container filled with water. The short-range correlations seen in Figure 9-9 are the result of the correlations in the detector outputs from correlated processes in the detector alone.

Wagner argued strongly against the simple engineering measures of contrast-to-noise ratio<sup>1</sup> (CNR) and pixel SNR<sup>2</sup> in his 1972 review because they are not task based. Furthermore, in terms of the random nature of noisy images, these measures do not capture the “color” or correlations in the noise (such as is seen in Figures 9-7 and 9-9), which can significantly impact an observer’s ability to perform simple visual tasks. The SNR<sub>t</sub> is both task based and dependent on the correlation properties of the noise.

Digital imaging systems have finite detector areas. Realistic data sets have interesting statistics, with mean pixel values that can depend on position because of gain variations in the detector or object structures, for example. In addition, the data may have correlations that vary with location in the image. Thus the assumptions for using the noise power spectrum to characterize the noise may not hold. One way to be sure is to check to see that the covariance matrix is indeed independent of such location-dependent variations. Another is to determine whether the full Fourier transformation of the covariance matrix, given by  $\mathcal{F} \mathbf{K} \mathcal{F}^\dagger$ , is indeed diagonal. A third approach is to compute a task SNR using both noise

<sup>1</sup> CNR = amplitude of an object divided by the standard deviation in the background.  
<sup>2</sup> pixel SNR = mean image at a specified location divided by the standard deviation at that location.



**Figure 9-9.** (a) Covariance matrix of water phantom. (b) Zoomed in image of covariance matrix in (a).

descriptions ( $\mathbf{K}$  and the NPS) to determine whether differences in the estimated system performance would result (Kyprianou et al. 2009; Brunner et al. 2010).

While the Fourier transformation may not yield a diagonal representation of the data covariance, it can be shown that there is *always* a transformation that will diagonalize the covariance matrix, giving

$$\mathbf{K} = \sum_{m=1}^M \mu_m \phi_m \phi_m^\dagger, \tag{9.26a}$$

or

$$\mathbf{K} = \Phi \mathbf{M} \Phi^\dagger. \tag{9.26b}$$

The  $\{\phi_n\}$  are yet another set of orthonormal basis functions, chosen in this case because of their unique ability to transform the covariance matrix into the diagonal matrix  $\mathbf{M}$ . The diagonalization of equation (9.26) is known as the Karhunen-Loève (KL) expansion. Where the noise power

spectrum gives the strength of the noise in each Fourier frequency, the KL expansion determines the strength of the fluctuations in the data at each basis vector  $\phi_n$ . Fourier analysis is simply the KL transformation in the special case of a wide-sense stationary random process.

Knowing that any set of basis functions can be used to represent the data, we can use these same KL vectors that diagonalize  $\mathbf{K}$  to represent any image, such that

$$\mathbf{g} = \sum_{m=1}^M \beta_m \phi_m. \quad (9.27)$$

We can make use of the discrete version of the general approach to finding coefficients of basis functions given in equation (9.16) to determine the  $\{\beta_m\}$ . We thus have two sets of basis functions that provide important information regarding the signal and noise properties of the imaging system, or the deterministic and stochastic properties, if you will. The SVD-derived basis functions tell us what information in the object is transferred (or not) from objects to images. The KL-derived basis functions describe the noise in the data in a set of independent channels.

#### 9.4.4 SNR Transfer for Digital Systems

Now that we have expressions for the data vector and its noise covariance at the output of a digital system, we can write the analogy to the ideal observer's test statistic in equation (9.14) as:

$$t(\mathbf{g}) = \Delta \bar{\mathbf{g}}^\dagger \mathbf{K}^{-1} \mathbf{g} = \Delta \bar{\boldsymbol{\beta}}^\dagger \mathbf{M}^{-1} \boldsymbol{\beta} = \sum_{m=1}^M \frac{\Delta \bar{\beta}_m \cdot \beta_m}{\mu_m}, \quad (9.28)$$

where  $\Delta \bar{\boldsymbol{\beta}}$  is the vector of coefficients describing the expected difference in the images between the two hypotheses, and  $\boldsymbol{\beta}$  is the vector of coefficients for the image being tested, as in equation (9.27). Equation (9.28) is the generalization of the prewhitening matched filter, showing that the optimal detector again correlates the data with the expected image of the signal, while again discounting those channels that contain the most noise.

The signal-to-noise ratio for the observer that uses the strategy in equation (9.28) is given by

$$\text{SNR}^2 = \Delta \bar{\mathbf{g}}^\dagger \mathbf{K}^{-1} \Delta \bar{\mathbf{g}} = \Delta \bar{\boldsymbol{\beta}}^\dagger \mathbf{M}^{-1} \Delta \bar{\boldsymbol{\beta}} = \sum_{m=1}^M \frac{(\Delta \bar{\beta}_m)^2}{\mu_m}. \quad (9.29)$$

The first expression for the  $\text{SNR}^2$ , in terms of  $\Delta \bar{\mathbf{g}}$  and  $\mathbf{K}^{-1}$ , is the observer's performance calculated using pixel means and covariances; hence we refer to it as the pixel-domain version of the figure of merit. We refer to the final expression, in terms of the coefficients of the KL basis, as the KL-domain figure of merit. For data made random by Poisson noise, which is independent and uncorrelated, the KL domain is the pixel domain; the covariance matrix of the pixel values is diagonal.

The signal-to-noise ratio presented by Bob Wagner in equation (9.1) is clearly evident in the general form of the  $\text{SNR}^2$  given in equation (9.29). Both are given in terms of coefficients that represent the signal to be detected in a data-space representation of independent, orthogonal channels that fully and exactly reproduce the signal. For LSIV systems, they are Fourier channels; for general digital systems they are the KL channels. In each case, the  $\text{SNR}^2$  gets its largest contributions from those channels or basis vectors for which the signal coefficients are large relative to the noise. In each case, the use of an orthogonal expansion results in a linear sum of contributions from each channel, making it easy to see how improvements in the signal transfer within a channel, or reductions in noise in a channel, impact the observer's performance.

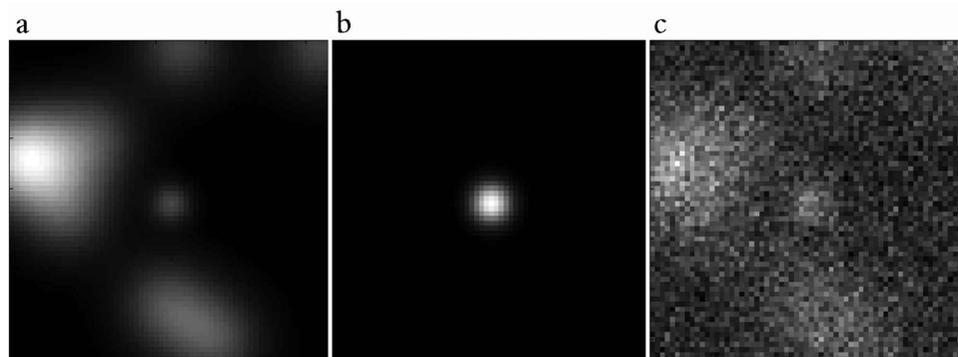
## 9.5 Current Trends

In this chapter we have presented a general form for the prewhitening matched filter for digital medical imaging systems. The Fourier version of this observer has been a familiar figure in the medical imaging literature since Bob Wagner's first papers on the subject. The pixel- or spatial-domain counterpart of this observer has received consideration attention as well, at least since the early 1990s, when it was suggested that the assumptions required for the Fourier version are not always valid, depending on the imaging modality and task.

Equation (9.29) gives the signal-to-noise ratio for the ideal observer when the data are additive and Gaussian, the same noise model considered by Wagner in his pioneering work. We must acknowledge that there are circumstances when a linear observer's performance can be surpassed by a nonlinear observer. This can happen whenever the data are not Gaussian distributed. The true ideal observer makes classification decisions based on the likelihood ratio. For non-Gaussian data, which can result from signal and/or background randomness, the ideal classifier may be a nonlinear function of the data.

The image to the left in Figure 9–10 is an example of nonGaussian image statistics. This image is generated by placing a random number of Gaussian blobs at random locations in the image, all of equal size and amplitude, to simulate a random, nonuniform background. In the center frame is an image of a single Gaussian object, the signal to be detected. On the right is a composite image containing a random lumpy background, the low-contrast Gaussian signal, and additive Gaussian noise (to simulate measurement noise). When a set of images is created with lumpy backgrounds like this example, a linear observer will not do as well as the optimal Bayesian observer, which computes a classifier that is a nonlinear function of the data.

If the data are not Gaussian, the linear observer of equation (9.28) can still be implemented, and that observer will



**Figure 9-10.** (a) Lumpy background. (b) Gaussian signal. (c) Noisy image of signal plus background.

still achieve an  $\text{SNR}^2$  given by equation (9.29). That is the best that can be achieved by an observer that uses a linear strategy for the classification task. In the medical imaging literature, the optimal linear observer is now commonly referred to as the Hotelling observer after Harold Hotelling, a statistician who published seminal papers in the early part of the 20<sup>th</sup> century on statistical hypothesis testing (Barrett et al. 1998). A number of studies have investigated the performance of the best linear observer for tasks in similar backgrounds to those in Figure 9–10, to better understand the limitations of the linear observer (Gallas and Barrett 2003; Park et al. 2007).

While the best linear observer may be suboptimal in overall detection performance, it has several selling points. As equation (9.28) tells us, its computation only requires knowledge of the expected data and the covariance of the data. These entities can be determined from a reasonably sized database for many applications. Or, these quantities may be determined via simulations of the imaging system under evaluation. Modern computers and accurate Monte Carlo algorithms for simulating gamma-ray and x-ray imaging physics make this a very viable option.

Another point in the linear observer's favor is that it has been shown to predict the performance of human observers for simple visual tasks. Wagner and Brown pointed out this selling feature in their 1985 paper. They noted at that time, though, that while the ideal observer's SNR tracks human performance in simple visual tasks, humans are relatively inefficient compared to the prewhitening matched filter for images containing negatively correlated noise, as is the case in CT. In such circumstances, the incorporation of a mechanism that models the eye-brain function is needed. Much work since 1985 has been done in the vision and psychophysics communities to validate simple modifications to the linear observer that model human performance mechanistically and give accurate predictions of human performance in practice. The most widely used such mechanism in the medical imaging community is the incorporation of finite-width channels, for example, octave-

wide channels in frequency space in the case of the Fourier-based ideal observer. Such channels impede the ability of the ideal observer to accurately prewhiten the images. Channelized linear observers have been shown to give a model for human performance that tracks human data for a wide range of image noise correlations, including the anticorrelated noise found in CT (Myers and Barrett 1987).

Channel mechanisms that model human performance can be used to explain certain inefficiencies of the human visual system. Channel mechanisms have a second use in image quality evaluation in medical imaging. *Efficient channels* are an important computational tool for allowing the estimation of the best linear observer's performance. By definition, efficient channels preserve the ideal observer's SNR, while reducing the computational burden of the image quality estimation problem through dimensionality reduction. Rather than dealing with an  $M$ -dimensional data set and an  $M \times M$ -dimensional covariance matrix, channels reduce the problem size to however many channels are used. An important problem of ongoing interest is the appropriate number and type of channel vectors to use for a particular task in medical imaging to enable accurate estimates of image quality.

Another of the current trends in applying the task-based SNR formalism to more realistic situations is to make use of the notion of efficient channels in cases where the ideal observer is a nonlinear function of the given data. The use of channels for this case is motivated by the large dimensionality problem that is a bottleneck to estimating the full nonlinear optimal observer's decision function and resulting performance. Images with nonGaussian lumpy backgrounds are a useful test set for comparing linear and nonlinear optimal observer strategies and performance results. Using this type of image, Park and her collaborators have investigated efficient channels for the ideal observer, including Laguerre-Gauss functions for extracting lesion characteristics (Park et al. 2007), singular vectors of a given linear imaging system (Park et al. 2009b), and weights from a partially least square algorithm that maximizes the covariance between the data

and the binary truth (Witten et al. 2010). These channels have been shown to be efficient to different degrees for both nonlinear and linear ideal observers.

## 9.6 Conclusion

Bob Wagner's earliest presentations and papers broke new ground in medical imaging, as he laid out a task-based framework for the assessment of image quality. His guiding principles are as important today as they were three decades ago. Medical imaging is under increased scrutiny in these modern times of comparative effectiveness research legislation and concerns regarding the overutilization of imaging tests. The answer to these questions and concerns is rooted in Bob's exhortation to evaluate the quality of a medical imaging system in terms of its usefulness in providing information that allows a specific user to perform a specific diagnostic task. In this chapter we have detailed the mathematical connections between the ideal observer SNR of Bob Wagner's pioneering work to figures of merit for modern digital imaging systems. Enormous progress is being made in the implementation of these methods using accurate simulations of objects and imaging systems, advanced numerical techniques, and experimental methods using relevant test objects in real systems.

## 9.7 Acknowledgment

This chapter is a distillation of concepts that can be found in greater detail in *Foundations of Image Science* by H. H. Barrett and K. J. Myers (2004). The authors thank R. M. Gagne for his development of Figures 9–2 through 9–6.

## 9.8 References

- Badano, A., I. S. Kyprianou, and J. Sempau. (2006). "Anisotropic imaging performance in indirect x-ray imaging detectors." *Med Phys* 33:2698–2713.
- Barrett, H. H., and K. J. Myers. *Foundations of Image Science*. Hoboken, NJ: John Wiley & Sons, Inc., 2004.
- Barrett, H. H., C. K. Abbey, B. D. Gallas, and M. Eckstein. (1998). "Stabilized estimates of Hotelling-observer detection performance in patient-structured noise." *Proc SPIE Medical Imaging* 3340:27–43.
- Brown, D. G. "Bob's first decade: In at the beginning," In: *Medical Imaging 2009: Image Perception, Observer Performance, and Technology Assessment*. B. Sahiner and D. J. Manning (eds.), Vol. 7263, 2009.
- Brunner, C. C., S. A. Hurowitz, S. A. Abboud, C. Hoeschen, and I. S. Kyprianou. "Noise characterization of computed tomography using the covariance matrix." *Proc SPIE Medical Imaging 2010: Physics of Medical Imaging*, Vol. 7622, 2010.
- Gallas, B. D., and H. H. Barrett. (2003). "Validating the use of channels to estimate the ideal linear observer." *J Opt Soc Am A* 20:1725–1738.
- Kyprianou, I. S., R. M. Gagne, A. Badano, S. Park, L. Thompson, B. D. Gallas, and K. J. Myers. "A practical method for measuring the H matrix of digital x-ray and cone beam CT imaging systems." *Proc SPIE Medical Imaging* 61421U1-12, 2006.
- Kyprianou, I. S., A. Badano, B. D. Gallas, and K. J. Myers. (2008). "Singular value description of a digital radiographic detector: theory and measurements." *Med Phys* 35(10):4744–4756.
- Kyprianou, I. S., A. Badano, S. Park, H. Liu, and K. J. Myers. "Noise and signal detection in digital x-ray detectors using the spatial definition of SNR" *Proc SPIE Medical Imaging*, Vol. 7258-44, 2009.
- Liu, H., I. S. Kyprianou, A. Badano, K. J. Myers, R. J. Jennings, S. Park, R. V. Kaczmarek, and K. Chakrabarti. "SKE/BKE task based ideal observer SNR methodology for mammography." *Proc SPIE Medical Imaging* 7258-44, 2009.
- Myers, K. J., and H. H. Barrett. (1987). "Addition of a channel mechanism to the ideal-observer model." *J Opt Soc Am A* 4:2447–2457.
- Park, S., H. H. Barrett, E. Clarkson, M. A. Kupinski, and K. J. Myers. (2007). "Channelized-ideal observer using Laguerre-Gauss channels in detection tasks involving non-Gaussian distributed lumpy backgrounds and a Gaussian signal." *J Opt Soc Am A* 24:B136–B150.
- Park, S., R. Leimbach, H. Liu, I. S. Kyprianou, R. J. Jennings, A. Badano, and K. J. Myers. "A task-based evaluation method for x-ray breast imaging systems using variable background phantoms." *Proc SPIE Medical Imaging* 7258-56, 2009a.
- Park, S., J. M. Witten, and K. J. Myers. (2009b). "Singular vectors of a linear imaging system as efficient channels for the Bayesian ideal observer." *IEEE Trans Med Imaging* 28(5): 657–668.
- Rao, N., M. E. Freed, and A. Badano. "Comparing experimental measurements of indirect x-ray detector responses with Monte Carlo predictions: Figures of merit and model improvement." *Proc SPIE Medical Imaging*, Vol. 7622, 2010.
- Rowe, R. K., J. N. Aarsvold, H. H. Barrett, J. C. Chen, W. P. Klein, B. A. Moore, I. W. Pang, D. D. Patton, and T. A. White. (1993). "A stationary hemispherical SPECT imager for three-dimensional brain imaging." *J Nucl Med* 34:474–480.
- Wagner, R. F., and K. E. Weaver. (1972). "An assortment of image quality indices for radiographic film-screen combinations—can they be resolved?" *Proc SPIE Appl Opt Instrum Med* 35:83–94.
- Wagner, R. F., and D. G. Brown. (1985). "Unified SNR analysis of medical imaging systems." *Phys Med Biol* 30(6):489–518.
- Witten, J. M., S. Park, and K. J. Myers. (2010). "Partial least squares: A method to compute efficient channels for the Bayesian ideal observers." *IEEE Trans Med Imaging* 29(4):1050–1058.

