Published in final edited form as:

Proc SPIE Int Soc Opt Eng. 2007 January 1; 8(34): 65150T.1–65150T.7. doi:10.1117/12.707804.

Adaptive Hotelling Discriminant Functions

Arthur Brème, Matthew A. Kupinski, Eric Clarkson, and Harrison H. Barrett

College of Optical Sciences Department of Radiology The University of Arizona, Tucson, AZ

Abstract

Any observer performing a detection task on an image produces a single number that represents the observer's confidence that a signal (e.g., a tumor) is present. A linear observer produces this test statistic using a linear template or a linear discriminant. The optimal linear discriminant is well-known to be the Hotelling observer and uses both first- and second-order statistics of the image data. There are many situations where it is advantageous to consider discriminant functions that adapt themselves to some characteristics of the data. In these situations, the linear template is itself a function of the data and, thus, the observer is nonlinear. In this paper, we present an example adaptive Hotelling discriminant and compare the performance of this observer to that of the Hotelling observer and the Bayesian ideal observer. The task is to detect a signal that is imbedded in one of a finite number of possible random backgrounds. Each random background is Gaussian but has different covariance properties. The observer uses the image data to determine which background type is present and then uses the template appropriate for that background. We show that the performance of this particular observer falls between that of Hotelling and ideal observers.

1. INTRODUCTION

Task-based assessment of image quality is increasingly utilized to evaluate medical imaging systems and reconstruction algorithms. One of the primary tasks in medical imaging is the detection of an abnormality such as a tumor. Any observer performing a tumor-detection task uses a discriminant function to map the high-dimensional image data to a single number, the test statistic. This test statistic is compared to a threshold to determine whether the observer decides the image is tumor absent or tumor present. Receiver operating characteristic (ROC) analysis^{1,2} has become the standard method for evaluating observers performing detection tasks. The area under the ROC curve (the AUC) is a common figure of merit used to assess signal-detection task performance.³ The observer with the highest possible AUC value is well-known to be the Bayesian ideal observer. This observer uses a discriminant function known as the likelihood ratio. Unfortunately, the likelihood ratio requires full knowledge of the distributions of the image data under the tumor-absent and tumor-present hypotheses. Even if these distributions are known, the computation of the likelihood ratio can still be cumbersome.⁴ An alternative approach is to use the ideal linear observer, i.e., the Hotelling observer. The Hotelling observer discriminant function uses

Corresponding author: M.A.K., kupinski@radiology.arizona.edu, Address: College of Optical Sciences, 1630 E. University Blvd., Tucson, AZ 85721.

only the first- and second-order statistics of the image data. Unfortunately, the Hotelling observer does not perform well in situations where the signal may be imbedded in a wide variety of background types. For example, a tumor may be located in different organs within the body and each tissue type produces background images with different statistical properties. In this case, lumping all of the tissues together in terms of their first- and second-order statistics degrades the performance of the Hotelling observer compared to a Hotelling observer that knows what tissue type is creating the background.

To address the problems noted above with the ideal and Hotelling observers, we are proposing an adaptive Hotelling observer. In general, an adaptive Hotelling observer first uses the image data to estimate the covariance of the background and then uses the corresponding Hotelling template to generate the test statistic. In this paper, the observer has prior knowledge about the types of backgrounds that may be present. Thus the observer first determines which of a finite number of background types is present and then uses the known covariance for that background to form the Hotelling test statistic. This observer is nonlinear because the form of the discriminant function depends on the image data. In practice, computing the full Hotelling discriminant function may be impractical since it involves the inversion of a very large matrix. In this paper, therefore, after the background type has been estimated, we use a channelized Hotelling observer to approximate the full Hotelling observer discriminant function.

The performance of the adaptive Hotelling observer should lie between that of the ordinary Hotelling observer and the Bayesian ideal observer. As such, it may offer a practical alternative for image-quality assessment when the ideal observer is too hard to compute and Hotelling observer performs poorly. In addition, the adaptive Hotelling discriminant may also be useful for modeling human observer performance on detection tasks.⁵ The adaptive Hotelling observer may also be useful in assessing adaptation rules for fully adaptive imaging systems.

2. ADAPTIVE HOTELLING DISCRIMINANTS

The ordinary Hotelling observer test statistic is given by,

$$t(\mathbf{g}) = \left[\bar{\mathbf{g}}_1 - \bar{\mathbf{g}}_0 \right]^\dagger K^{-1} \left[\mathbf{g} - \bar{\mathbf{g}}_0 \right] - \frac{1}{2} \left[\bar{\mathbf{g}}_1 - \bar{\mathbf{g}}_0 \right]^\dagger K^{-1} \left[\bar{\mathbf{g}}_1 - \bar{\mathbf{g}}_0 \right] \quad (1)$$

where $\mathbf{g}$ is the image data, $\bar{\mathbf{g}}_1$ is the mean image data under the signal-present hypothesis (H_1), $\bar{\mathbf{g}}_0$ is the mean image data under the signal-absent hypothesis (H_0), and K is the image-data covariance matrix. This equation differs from the standard form for the Hotelling test statistic since the standard form usually drops all constant terms. These additional terms must be retained for this work as will be discussed later. A simpler form for Eqn. 1 is,

$$t(\mathbf{g}) = \omega^\dagger \mathbf{g} + c, \quad (2)$$

where $\omega = K^{-1} \left[\bar{\mathbf{g}}_1 - \bar{\mathbf{g}}_0 \right]$, and c is the constant term given by,

$$c = - \left[\bar{\mathbf{g}}_1 - \bar{\mathbf{g}}_0 \right]^\dagger K^{-1} \bar{\mathbf{g}}_0 - \frac{1}{2} \left[\bar{\mathbf{g}}_1 - \bar{\mathbf{g}}_0 \right]^\dagger K^{-1} \left[\bar{\mathbf{g}}_1 - \bar{\mathbf{g}}_0 \right]. \quad (3)$$

The general form for the adaptive Hotelling observer is similar to Eqn. 2 and is given by,

$$t(\mathbf{g}) = \omega(\mathbf{g})^\dagger \mathbf{g} + c(\mathbf{g}), \quad (4)$$

where both the template $\mathbf{w}$ and the constant term c now depend on the image data itself. Because the template and the constant depend on the image data $\mathbf{g}$, this observer is a nonlinear observer. Specifically, the adaptive Hotelling observer computes a test statistic using a two-stage process. First, the adaptive observer analyzes the image data $\mathbf{g}$ and determines which type of background, out of a finite number of classes, is present in the image. After this background classification step, the observer then uses the appropriate template $\mathbf{w}$ and the constant term c for that background type to produce a test statistic. The performance of this observer depends on the performance of both the background classification step and the Hotelling observer step. The constant term $c(\mathbf{g})$ is needed for this observer to ensure that the test statistics for different background types are comparable.

3. SIMULATION

To evaluate the adaptive Hotelling observer, we conducted a simulation study designed so that the performance of this observer could be compared to the performance of both the ideal and the Hotelling observers.

3.1. Image generation

For this initial study evaluating the adaptive Hotelling observer, we employed simulated images that have both a random background component to simulate object variability and a white noise component to simulate detector noise. Images were generated using one of five different types of backgrounds. These backgrounds were produced by convolving white Gaussian noise images with filters (denoted by I_i) of varying widths. The images were 128×128 pixels and the filters were Gaussians with standard deviations (i.e., widths) ranging from 4.48 to 22.4 pixels. A constant background of 120 was added to the filtered images and we then added Gaussian white noise with a standard deviation of 5. Example tumor-absent images are shown in Fig. 1 for the 5 types of backgrounds. Because the overall distribution of the images is Gaussian, we can compute the performance of the Bayesian ideal observer and compare this to the performance of the adaptive Hotelling observer.

For all background types, the signal to be detected was a circular Gaussian function with a standard deviation of 0.16 and amplitude 8. Figure 2 shows the signal to be detected without any background at all. Note that the gray scales in Figures 2 and 1 are normalized to the maximum and minimum values in the images and, thus, cannot be compared to one another. Visual inspection of the signal imbedded in all 5 types of backgrounds indicates that the signal is very difficult for humans to detect.

3.2. Background classification

The first stage for the adaptive Hotelling observer is the background classification stage. The goal is to determine which of the five possible types of backgrounds is present. If we represent the full covariance matrices for the five background types with the term K_i , where i goes from 1 to 5, then the likelihood for the image data conditioned on the background i hypothesis (B_i), is given by,

$$pr(\mathbf{g}|B_i) = \frac{1}{\sqrt{(2\pi)^M \det(K_i)}} \exp \left[-\frac{1}{2} (\mathbf{g} - \bar{\mathbf{g}})^\dagger K_i^{-1} (\mathbf{g} - \bar{\mathbf{g}}) \right], \quad (5)$$

where M is the number of pixels in the image data and $\bar{\mathbf{g}}$ is the mean image data under the signal-absent hypothesis. The lack of a subscript on the $\bar{\mathbf{g}}$ is due to the fact that the mean image data was set to be the same for all 5 types of backgrounds. Furthermore, the signal is weak and is ignored in the mean data for the purposes of background classification. The mean signal is important for the second step of the adaptive Hotelling observer. The likelihood (Eqn. 5) is computed for all 5 background types for an image $\mathbf{g}$, and the background type with the highest likelihood is selected.

In practice, computing the likelihood in Eqn. 5 is difficult since it requires the inversion of a large matrix. To alleviate this issue we make the poor assumption that the background is stationary and the covariance matrix can be diagonalized using a discrete Fourier transform. While the stationarity is, in general, a poor assumption, it does provide us with a rapid method for classifying the background type. In the Results section, we will study the performance of the adaptive Hotelling observer as the performance of the background classification stage varies. Under the assumption of stationarity, Eqn. 5 simplifies to,

$$pr(\mathbf{g}|B_i) = \frac{1}{\sqrt{(2\pi)^M \prod_{m=1}^M [\Gamma_i]_m}} \exp \left[-\frac{1}{2} \mathbf{G}^\dagger \mathbf{G} (\text{Diag}(\Gamma_i))^{-1} \right], \quad (6)$$

where $\mathbf{G}$ is the discrete Fourier transform (DFT) of $\mathbf{g} - \bar{\mathbf{g}}$ and Γ_i the DFT of the i th filter $\mathbf{l}_i$ squared plus the white-noise variance, i.e.,

$$\Gamma_i = \mathcal{F}\{\mathbf{l}_i\}^2 + \sigma^2. \quad (7)$$

The matrix $\text{Diag}(\Gamma_i)$ is diagonal with components given by the vector Γ_i .

3.3. Channelized Hotelling Observer

The second stage of the adaptive Hotelling observer is to use Hotelling template for the estimated background type to compute a test statistic. In practice, computing the full Hotelling observer is burdensome because the full Hotelling observer requires the inversion of a large matrix. Thus, we employ a channelized Hotelling observer using Laguerre-Gauss channels which have been shown to approximate the performance of the full Hotelling observer.⁶ The channelized Hotelling observer uses a channel operator T to lower the dimension of the data vector $\mathbf{g}$. For a given image $\mathbf{g}$ the channel outputs are given by,

$$v = Tg \quad (8)$$

where T is the $C \times M$ matrix of the Laguerre-Gauss channels. The test statistic using the channelized Hotelling observer becomes,

$$t(g) = \left[\hat{v}_1 - \hat{v}_0 \right]^\dagger \left[K_v^{(k)} \right]^{-1} \left[v - \hat{v}_0 \right] - \frac{1}{2} \left[\hat{v}_1 - \hat{v}_0 \right]^\dagger \left[K_v^{(k)} \right]^{-1} \left[\hat{v}_1 - \hat{v}_0 \right] \quad (9)$$

where hats denote estimates from samples. The term $\left[K_v^{(k)} \right]$ is the estimated channel-covariance matrix for the k th background type. The mean channel vectors $\hat{v}_i$ do not depend on background type since the mean image data vector is a constant for all background types.

For our study, the number of Laguerre-Gauss channels used was $C = 5$. A total of 500

images for each background type were used to estimate the covariance matrices $\left[K_v^{(k)} \right]$. As stated earlier, the extra terms in Eqn. 9 ensure that test statistics for different background types can be compared to one another.

4. RESULTS

To evaluate the adaptive Hotelling observer, we simulated 5,000 image pairs. The probability $P(B_i)$ for each background type was 1/5. First we evaluated the background-classification stage with these 5,000 images and found that the background classifier correctly identified the background 98.6 % of time. To study the performance of the second stage of the adaptive Hotelling observer, we first analyzed the distributions of the test statistics for the various background types assuming the first stage correctly classified the background. Figure 3 shows the distributions for the images of the different background types. Note that the separation between the classes for the various background types differs greatly. Further note the symmetry around 0 between the two classes of data. This symmetry is predicted by our form of the Hotelling observer shown in Eqn. 1.

We now compare the overall performance of the adaptive Hotelling observer (AHO) to that of the standard Hotelling observer (HO) and the Bayesian ideal observer (BIO). The performance of the HO was estimated using the channelized Hotelling observer with Laguerre-Gauss channels. Unlike the AHO, the covariance matrix K_v employed was an average over all background types. Because our image data are a Gaussian mixture model, we can compute the performance of the BIO. The likelihood ratio for the BIO is given by,

$$t_{IO} = \frac{\sum_{i=1}^5 pr(g|B_i, H_1) pr(B_i)}{\sum_{i=1}^5 pr(g|B_i, H_0) pr(B_i)} \quad (10)$$

where,

$$pr(g|B_i, H_j) = \frac{1}{\sqrt{(2\pi)^M \prod_{m=1}^M [\Gamma_i]_m}} \exp \left[-\frac{1}{2} \mathbf{G}_j^\dagger \mathbf{G}_j (\text{Diag}(\Gamma_i))^{-1} \right]. \quad (11)$$

In Eqn. 11, the term $\mathbf{G}_j$ is the Fourier transform of $\mathbf{g} - \bar{\mathbf{g}}_j$. Note that this differs from the form of $\mathbf{G}$ used in Eqn. 6. The ROC curves for both the AHO, the HO, and the BIO are shown in Fig. 4. The performance of the adaptive Hotelling observer falls between that of the Bayesian ideal observer and the Hotelling observer. The AUCs for the BIO, the HO, and the AHO are 0.893 ± 0.004 , 0.672 ± 0.005 , and 0.886 ± 0.003 , respectively. The error rate for the AUCs were determined using standard MRMC methods.⁷

The performance of the AHO is close to that of the BIO thanks to the good performance of the background classification stage. To fully assess the performance of the AHO under all conditions, we voluntarily decreased the performance of this first stage. We examined two situations: one in which the background-classification stage randomly selected one of the 5 backgrounds, and the other in which the background-classification stage was operating at 84% accuracy. For 80% of the images, the actual background of the image was selected, and for the rest, one background type was selected randomly (a fifth of these will be correct which accounts for the additional 4%). We see in Fig. 5 that the performance of the AHO drops below the BIO in both situations. We further find that the lowest performance of the AHO equals that of the HO (see Fig. 5(a)). Thus we hypothesize that the AUC of the AHO has a lower bound near that of the AUC for the HO.

5. CONCLUSIONS AND FUTURE WORK

In this paper, we presented an adaptive Hotelling observer that first selects a background type and then uses that information to compute the test statistic for the Hotelling observer. We used a very simple background-classification algorithm based on stationarity assumptions. In practice, the algorithm used to classify the background may need to be different. We compared the performances of the adaptive Hotelling observer to that of the ideal observer and the Hotelling observer. We found that the adaptive Hotelling observer is easy to compute and performs better than the standard Hotelling observer when the background type varies from image to image. In the future, we hope to compare the performance of the adaptive Hotelling observer to that of human observers. It has been suggested that humans adapt to the background when performing detection tasks.⁵

To reduce the computational burden of the Hotelling observer, we reduced the dimension of the image data by a projection on channels. To simplify the background classification, we made the poor assumption that the background was stationary. One of our future goals will be to try to channelize the background classification, instead of using the stationary assumption. Finally, we are currently developing an adaptive SPECT, small-animal imaging system that adapts the geometry of the system based on task performance and initial scans of the specific animal. We hope to use our work on the adaptive Hotelling observer to help devise figures of merit for rapidly optimizing the adaptive imaging system.

6. ACKNOWLEDGMENTS

This work was supported under NIH/NIBIB grants R01-EB002146, R37 EB000803, and P41 EB002035.

REFERENCES

1. Metz CE. Basic principles of ROC analysis. *Seminars in Nuclear Medicine* VIII. 1978;283–298.
2. Swets J. ROC analysis applied to the evaluation of medical imaging techniques. *Investigative Radiology*. 1979; 14:109–121. [PubMed: 478799]
3. Hanley JA, McNeil BJ. The meaning and use of the area under a receiver operating characteristic (ROC) curve. *Radiology*. 1982; 143:29–36. [PubMed: 7063747]
4. Kupinski MA, Hoppin JW, Clarkson E, Barrett HH. Ideal-observer computation in medical imaging with use of Markov-chain Monte Carlo techniques. *JOSA A*. 2003; 20(3):430–438. [PubMed: 12630829]
5. Zhang Y, Abbey CK, Eckstein MP. Adaptive detection mechanisms in globally statistically nonstationary-oriented noise. *Journal of the Optical Society of America A*. 2006; 23(7):1549–1558.
6. Barrett, HH.; Abbey, C.; Gallas, B.; Eckstein, M. Stabilized estimates of hotelling-observer detection performance in patient-structured noise. In: Kundel, HL., editor. *Medical Imaging 1998: Image Perception*. Vol. 3340. SPIE; 1998. p. 27–43.
7. Clarkson E, Kupinski MA, Barrett HH. A probabilistic development of the MRMC method. *Academic Radiology*. 2006; 13(10)

Figure 1.
Samples from the five different background ensembles, (a) $i = 1$, to (e) $i = 5$.

Figure 2.
Image of the signal to be detected.

Figure 3.

The distributions of the test statistics for the second stage of the adaptive Hotelling observer. These distributions were generated assuming that the background-classification stage is operating at 100% performance. (a) type $i = 1$, to (e) type $i = 5$.

Figure 4.

The ROC curves for the three observer studied in this paper. The ROC curve for the AHO, as predicted, falls between that of BIO and the HO. The performance of the BIO and the AHO are very similar because the performance of the background-classification stage was at 98.6%.

Figure 5.

Comparison of the ROC curve for the 3 observers. In (a) the background-detection is replaced by a guessing observer. In (b), the background-classification stage was operating at an accuracy of 84%. The AUCs for the AHO are (a) 0.683 ± 0.005 and (b) 0.827 ± 0.004 .