Published in final edited form as:

J Opt Soc Am A Opt Image Sci Vis. 2020 February 01; 37(2): 174–181. doi:10.1364/JOSAA.37.000174.

Probability of error for detecting a change in a parameter and Bayesian Fisher information

Eric Clarkson

Department of Medical Imaging and College of Optical Sciences University of Arizona, Tucson, Arizona 85721, USA

Abstract

The van Trees inequality relates the ensemble mean squared error of an estimator to a Bayesian version of the Fisher information. The Ziv-Zakai inequality relates the ensemble mean squared error of an estimator to the minimum probability of error for the task of detecting a change in the parameter. In this work we complete this circle by deriving an inequality that relates this minimum probability of error to the Bayesian version of the Fisher information. We discuss this result for both scalar and vector parameters. In the process we discover that an important intermediary in the calculation is the total variation of the posterior probability distribution function for the parameter given the data. This total variation is of interest in its own right since it may be easier to compute than the other figures of merit discussed here.

1. INTRODUCTION

In imaging applications we often use the data generated by the imaging system to perform a task of interest. Such a task may be an estimation task, where we want to estimate a set of parameters associated with the object being imaged; a classification task, where we want to decide to which class an object belongs; or a task that combines estimation with classification. For example, in medical imaging we may be interested in estimating the size and shape of a tumor to determine whether the prescribed therapy is having a positive effect. A closely related classification task is to determine if the parameters defining the size and shape of the tumor have changed or remained the same. In remote sensing we may want to estimate the location and orientation of a particular target object. A closely related classification task is to determine whether the target object has moved or not.

The Fisher information (FI) and the Fisher information matrix (FIM) are fundamental concepts in statistical estimation theory. For a scalar parameter the well-known Cramer–Rao bound (CRB) shows that the inverse of the FI is a lower bound for the variance of an unbiased estimator of the parameter. For reference we define the FI and state the CRB for a scalar parameter in Section 2. For a vector parameter the inverse of the FIM provides a similar lower bound for the covariance matrix of an unbiased estimator. Less well known is the connection between the FI and the FIM to signal detection theory. For a scalar parameter

clarkson@radiology.arizona.edu.

Disclosures. The author declares that there are no conflicts of interest related to this paper.

we can ask how well we can detect a small change in that parameter from noisy data associated with it via a conditional probability distribution function (PDF). The optimal method for detecting such a change is to compute the likelihood ratio and compare it to a threshold. Such an observer is called an ideal observer and, by varying the threshold, we can plot the receiver operating characteristic (ROC) curve for the ideal observer. The area under this curve (AUC), the ideal-observer AUC, is a figure of merit measuring the quality of the data for the task of detecting the change in the parameter. For small parameter changes, the ideal-observer AUC is, to first order, proportional to the FI. This connection between the FI and our ability to detect a small change in a parameter is reviewed briefly in Section 3. Our main goal in this work is to find a similar connection between the Bayesian version of the FI and this detection task. For vector parameters the connection between the ideal-observer AUC for the detection of a small change in the parameter vector and the FIM is described in Section 4. We will also find a similar connection between this detection task and the Bayesian FIM. The ideal-observer AUC and the FIM do not depend on the prior probabilities for the signal present and signal absent hypotheses. The new element in the Bayesian approach is a prior on the parameters governing the statistics of the data, which can be used to define these prior probabilities. The resulting minimum probability of error (MPE) becomes the quantity of interest for the detection task that we will relate to the Bayesian FIM.

In Section 5 we introduce the Ziv-Zakai inequality in our notation. The detection task relevant to this inequality is the detection task we will be considering in the subsequent sections. In this task we are trying to detect a change in a parameter but we have more information than we do in the task described in Section 3. In particular, we have a prior distribution on the parameter and use this to define prior probabilities for the two parameter values that represent the two hypotheses in the detection task. The Ziv-Zakai inequality relates the MPE for this task to the ensemble mean squared error (EMSE) for any estimator of the parameter. The MPE is the probability of error for the ideal observer in the detection task using an optimal threshold determined by the prior probabilities of the two hypotheses. In Section 6 we briefly review the van Trees inequality, which relates this EMSE to the Bayesian FI. These two sections provide context for Sections 7 and 8, which include a statement of the main results of this paper and a derivation of these results: relations between the MPE for the Ziv-Zakai detection task, the total variation (TV) of the posterior PDF, and the Bayesian FI. The posterior TV may be a useful figure of merit in its own right since it gives us the first-order approximation of the MPE when the two parameter values are close to each other. The vector-parameter version of this result is given in Section 9.

If the posterior PDF is unimodal, then the posterior TV is easy to compute, as shown in Section 10. We compute the posterior TV and Bayesian FI for two examples of unimodal posterior PDFs in this section. We then compute the posterior PDFs, posterior TVs, and Bayesian FIs for two imaging applications, one with a scalar parameter and one with a vector parameter. Finally, in the Conclusion we summarize our results and their implications for the evaluation of the performance of measurement systems on detection and estimation tasks.

2. FISHER INFORMATION

For most of this paper we will be using a scalar parameter θ and a conditional PDF $\text{pr}(\mathbf{g}|\theta)$ for the data vector $\mathbf{g}$. This data vector may, for example, be the end result of an imaging experiment. However, all of the results generalize to a vector parameter $\boldsymbol{\theta}$ and we will indicate those generalizations as we proceed. Before getting to the main new results, and to establish notation, we first review some concepts relevant to estimation tasks, detection tasks, and the connections between them. In all that follows angle brackets indicate the probabilistic expectation and the subscripts on the angle brackets indicate which random variables or vectors are averaged over and, if needed, which variables or vectors are held fixed. For example, the subscript $\mathbf{g}|\theta$ means that we are using the conditional PDF $\text{pr}(\mathbf{g}|\theta)$ to average over $\mathbf{g}$ with θ held fixed.

The FI for the parameter of interest θ is given by the expectation [1]

$$F(\theta) = \left\langle \left[\frac{d}{d\theta} \ln \text{pr}(\mathbf{g}|\theta) \right]^2 \right\rangle_{\mathbf{g}|\theta}. \quad (1)$$

If $\hat{\theta}(\mathbf{g})$ is an estimator of θ from the data, then this estimator is unbiased if

$$\langle \hat{\theta}(\mathbf{g}) \rangle_{\mathbf{g}|\theta} = \theta.$$

The CRB then states that, given certain conditions on the conditional PDF, the variance of any unbiased estimator satisfies [1]

$$\text{var}(\hat{\theta}) \geq \frac{1}{F(\theta)}. \quad (2)$$

Thus the FI is an important quantity when we are considering estimation tasks. In the next section we discuss a not-so-well-known relation between the FI and a specific detection task.

3. FI AND IDEAL-OBSERVER AUC

In this section we introduce a specific binary classification task that is related to the estimation task in the previous section and that we will be considering throughout this paper. We suppose that we are given the data vector $\mathbf{g}$ and told that one of two hypotheses is true. The hypothesis H_0 is that $\mathbf{g}$ is a sample drawn from the PDF $\text{pr}(\mathbf{g}|\theta)$, which we write as $\mathbf{g} \sim \text{pr}(\mathbf{g}|\theta)$. The hypothesis H_1 is that $\mathbf{g} \sim \text{pr}(\mathbf{g}|\tilde{\theta})$. The job of the observer is to determine which hypothesis is true. The optimal observer for this task by many metrics, some of which we will be discussing below, is the Bayesian ideal observer, also known simply as the ideal observer [2,3]. This observer computes the ratio of the H_1 PDF to the H_0 PDF, which is called the likelihood ratio, and which we will denote by $\Lambda(\mathbf{g}|\theta, \tilde{\theta})$. The ideal observer then compares the likelihood ratio to a threshold that we will call y . If $\Lambda(\mathbf{g}|\theta, \tilde{\theta}) > y$, then the ideal observer concludes that $\mathbf{g} \sim \text{pr}(\mathbf{g}|\tilde{\theta})$, i.e., that hypothesis H_1 is true. Otherwise, this observer declares that $\mathbf{g} \sim \text{pr}(\mathbf{g}|\theta)$ and hypothesis H_0 is true.

Due to noise in the data vector quantified by the PDFs $\text{pr}(\mathbf{g}|\theta)$ and $\text{pr}(\mathbf{g}|\tilde{\theta})$, the ideal observer, although optimal, is not always right. One possible error is a false positive (FP) where $\mathbf{g} \sim \text{pr}(\mathbf{g}|\theta)$ but $\Lambda(\mathbf{g}|\theta, \tilde{\theta}) > y$. Note that when $\mathbf{g} \sim \text{pr}(\mathbf{g}|\theta)$ the likelihood ratio $\Lambda = \Lambda(\mathbf{g}|\theta, \tilde{\theta})$ is a random variable with a PDF that we will denote by $\text{pr}_0(\Lambda|\theta, \tilde{\theta})$. The subscript indicates the hypothesis that is in force, and the θ and $\tilde{\theta}$ after the vertical bar are there because the function $\Lambda(\mathbf{g}|\theta, \tilde{\theta})$, and hence the PDF for Λ , depends on both of these variables. The probability of an FP outcome is known as the false positive fraction (FPF). The FPF will be written as $\text{FPF}(y|\theta, \tilde{\theta})$. The other possible error is a false negative (FN) where $\mathbf{g} \sim \text{pr}(\mathbf{g}|\tilde{\theta})$ but $\Lambda(\mathbf{g}|\theta, \tilde{\theta}) \leq y$. We will write the false negative fraction (FNF) as $\text{FNF}(y|\theta, \tilde{\theta})$.

The true positive fraction (TPF) is defined by $\text{TPF}(y|\theta, \tilde{\theta}) = 1 - \text{FNF}(y|\theta, \tilde{\theta})$ and is the probability that $\Lambda(\mathbf{g}|\theta, \tilde{\theta}) > y$ when H_1 is valid. For a given pair $(\theta, \tilde{\theta})$, the corresponding ROC curve is a plot of $\text{TPF}(y|\theta, \tilde{\theta})$ versus $\text{FPF}(y|\theta, \tilde{\theta})$. The area under this curve is the ideal-observer AUC, and it can be used as a figure of merit for the quality of the data with respect to the classification task. One advantage of the AUC as a figure of merit is that it can be estimated from a two alternative forced choice test without actually plotting the ROC curve. We will use the notation $\text{AUC}(\theta, \tilde{\theta})$ for this area. The ideal-observer detectability $d(\theta, \tilde{\theta})$ is an alternative to $\text{AUC}(\theta, \tilde{\theta})$ as a figure of merit and is defined by

$$\text{AUC}(\theta, \tilde{\theta}) = \frac{1}{2} + \frac{1}{2} \text{erf} \left[\frac{1}{2} d(\theta, \tilde{\theta}) \right]. \quad (3)$$

The ideal-observer AUC always satisfies $0.5 \leq \text{AUC}(\theta, \tilde{\theta}) \leq 1$ and the ideal-observer detectability is always a non-negative real number.

In previous work it was shown that the ideal-observer detectability is related to the FI in a Taylor series expansion as [4–6]

$$d^2(\theta, \theta + \Delta\theta) = F(\theta)(\Delta\theta)^2 + \dots \quad (4)$$

Thus, the FI has an interpretation in terms of detecting a small change in a parameter. The next term in the Taylor series, the cubic term, is known and involves the derivative of the FI. Thus, it is possible to estimate the error in the second-order approximation to the detectability. In order to compare this result to some results derived below we want to relate the FI directly to $\text{AUC}(\theta, \theta + \Delta\theta)$. To do this we start with the Taylor series for the error function:

$$\text{erf}(z) = \frac{2}{\sqrt{\pi}} \left(z - \frac{z^3}{3} + \dots \right). \quad (5)$$

Using this series we find that

$$\text{AUC}(\theta, \theta + \Delta\theta) = \frac{1}{2} + \frac{1}{2\sqrt{\pi}} |\Delta\theta| \sqrt{F(\theta)} + \dots \quad (6)$$

Note that the absolute value here indicates that $\text{AUC}(\theta, \tilde{\theta})$ is not differentiable with respect to $\tilde{\theta}$ at $\tilde{\theta} = \theta$. We can, however, formulate this equation in terms of one-sided derivatives as

$$\left. \frac{d}{d\tilde{\theta}} \text{AUC}(\theta, \tilde{\theta}) \right|_{\tilde{\theta} = \theta \pm} = \pm \frac{1}{2\sqrt{\pi}} \sqrt{F(\theta)}. \quad (7)$$

This formulation can be compared to results below relating the minimum probability of error on this classification task to the Bayesian version of the FI.

4. VECTOR VERSION OF FI AND AUC

For a p -dimensional vector parameter θ the $p \times p$ FIM is defined by [1]

$$\mathbf{F}(\theta) = \left\langle [\nabla_{\theta} \ln \text{pr}(\mathbf{g}|\theta)] [\nabla_{\theta} \ln \text{pr}(\mathbf{g}|\theta)]^{\dagger} \right\rangle_{\mathbf{g}|\theta}. \quad (8)$$

This matrix is related to the ideal-observer detectability of a small change in the parameter vector by the Taylor series expansion [4–6]

$$d^2(\theta, \theta + \Delta\theta) = \Delta\theta^{\dagger} \mathbf{F}(\theta) \Delta\theta + \dots \quad (9)$$

If $\mathbf{u}$ is an arbitrary unit vector in parameter space, then we have the one-sided derivatives d

$$\left. \frac{d}{dt} \text{AUC}(\theta, \theta + t\mathbf{u}) \right|_{t=0+} = \frac{1}{2\sqrt{\pi}} \sqrt{\mathbf{u}^{\dagger} \mathbf{F}(\theta) \mathbf{u}} \quad (10)$$

and

$$\left. \frac{d}{dt} \text{AUC}(\theta, \theta + t\mathbf{u}) \right|_{t=0-} = -\frac{1}{2\sqrt{\pi}} \sqrt{\mathbf{u}^{\dagger} \mathbf{F}(\theta) \mathbf{u}}. \quad (11)$$

For example, when $p = 2$ we can plot $\text{AUC}(\theta, \tilde{\theta})$ as a function of $\tilde{\theta}$ for a fixed θ . This surface will descend to a point at $\tilde{\theta} = \theta$. The slope as we descend to or ascend from this singularity in the direction $\mathbf{u}$ is determined by the quantity $\mathbf{u}^{\dagger} \mathbf{F}(\theta) \mathbf{u}$. A larger value for this slope implies that it will be easier to detect a small change in the parameter vector in that direction from the data that we have to work with.

5. THE ZIV-ZAKAI INEQUALITY

The setting for the results we will be describing below relating the MPE in the classification task described above to the Bayesian FI is the same as the setting for the Ziv-Zakai inequality that relates the EMSE of an estimator to this same MPE. We will briefly discuss this inequality in order to introduce this setting and some notation we will be using. If we have a prior PDF $\text{pr}(\theta)$ on the parameter of interest, then we may define probabilities for the two hypotheses H_0 and H_1 via

$$\Pr_0(\theta, \tilde{\theta}) = \frac{\text{pr}(\theta)}{\text{pr}(\theta) + \text{pr}(\tilde{\theta})} \quad (12)$$

and

$$\Pr_1(\theta, \tilde{\theta}) = \frac{\text{pr}(\tilde{\theta})}{\text{pr}(\theta) + \text{pr}(\tilde{\theta})}. \quad (13)$$

To minimize the probability of error the optimal threshold is given by

$$y = y(\theta, \tilde{\theta}) = \frac{\Pr_0(\theta, \tilde{\theta})}{\Pr_1(\theta, \tilde{\theta})} = \frac{\text{pr}(\theta)}{\text{pr}(\tilde{\theta})}. \quad (14)$$

The MPE in this setting is therefore

$$P_e(\theta, \tilde{\theta}) = \Pr_0(\theta, \tilde{\theta}) \int_{y(\theta, \tilde{\theta})}^{\infty} \text{pr}_0(\lambda|\theta, \tilde{\theta}) d\lambda + \Pr_1(\theta, \tilde{\theta}) \times \int_0^{y(\theta, \tilde{\theta})} \text{pr}_1(\lambda|\theta, \tilde{\theta}) d\lambda. \quad (15)$$

The two terms here correspond to the FP and FN cases. If an observer simply decides H_0 is true every time, then the probability of error is $\Pr_1(\theta, \tilde{\theta})$. Similarly if an observer simply decides H_1 is true every time, then the probability of error is $\Pr_0(\theta, \tilde{\theta})$. Therefore, for the MPE we always have $0 \leq P_e(\theta, \tilde{\theta}) \leq \min\{\Pr_0(\theta, \tilde{\theta}), \Pr_1(\theta, \tilde{\theta})\}$. Since $\Pr_0(\theta, \tilde{\theta}) + \Pr_1(\theta, \tilde{\theta}) = 1$, we can further conclude that $0 \leq P_e(\theta, \tilde{\theta}) \leq 0.5$.

The EMSE for an estimator $\hat{\theta}(\mathbf{g})$ is given by

$$\text{EMSE}(\hat{\theta}) = \left\langle \left[\hat{\theta}(\mathbf{g}) - \theta \right]^2 \right\rangle_{\mathbf{g}|\theta}. \quad (16)$$

The Ziv-Zakai inequality can now be written as [7–9]

$$\text{EMSE}(\hat{\theta}) \geq \frac{1}{2} \left\langle \int_{-\infty}^{\infty} P_e(\theta, \tilde{\theta}) |\tilde{\theta} - \theta| d\tilde{\theta} \right\rangle_{\theta}. \quad (17)$$

A large value of $P_e(\theta, \tilde{\theta})$ when $|\tilde{\theta} - \theta|$ is small is intuitively expected, as is a small value of $P_e(\theta, \tilde{\theta})$ when $|\tilde{\theta} - \theta|$ is large. The Ziv-Zakai inequality shows that a large value of $P_e(\theta, \tilde{\theta})$ when $|\tilde{\theta} - \theta|$ is also large is very bad for the estimation problem as it will force a large EMSE for any estimator. We will be showing below that the behavior of $P_e(\theta, \tilde{\theta})$ when $|\tilde{\theta} - \theta|$ is small is related to the total variation of the posterior PDF for θ , and to the Bayesian FI.

6. BAYESIAN FI AND EMSE

What we have been referring to as the Bayesian FI is given by

$$F = \langle F(\theta) \rangle_{\theta} + \left\langle \left[\frac{d}{d\theta} \ln \text{pr}(\theta) \right]^2 \right\rangle_{\theta}. \quad (18)$$

The posterior PDF $\text{pr}(\theta|\mathbf{g})$ for θ is defined by the equation $\text{pr}(\theta|\mathbf{g})\text{pr}(\mathbf{g}) = \text{pr}(\mathbf{g}|\theta)\text{pr}(\theta)$, where

$$\text{pr}(\mathbf{g}) = \int_{-\infty}^{\infty} \text{pr}(\mathbf{g}|\theta)\text{pr}(\theta)d\theta.$$

In terms of this posterior PDF the Bayesian FI can also be written as

$$F = \left\langle \left\langle \left[\frac{d}{d\theta} \ln \text{pr}(\theta|\mathbf{g}) \right]^2 \right\rangle_{\mathbf{g}|\theta} \right\rangle_{\theta}. \quad (19)$$

This expression will be useful when we relate $P_e(\theta, \tilde{\theta})$ to the Bayesian FI. The usual application of the Bayesian FI is the van Trees inequality, also called the Bayesian CRB, which states that [10,11]

$$\text{EMSE}(\hat{\theta}) \geq F^{-1}. \quad (20)$$

There are versions of the Ziv-Zakai inequality and the van Trees inequality for vector parameters but we will not be discussing those here. We would be completing the circle started by the Ziv-Zakai and van Trees inequalities if we had a relation between $P_e(\theta, \tilde{\theta})$ and F . This is the subject of the next section and we will find that an intermediary in this relation is the total variation of the posterior PDF $\text{pr}(\theta|\mathbf{g})$.

7. BAYESIAN FI AND MPE

The main results of this paper, which are derived in the next section, involve one-sided derivatives with respect to $\tilde{\theta}$ of the $\text{MPE}P_e(\theta, \tilde{\theta})$ as we approach $\tilde{\theta} = \theta$. In the first relation we use the prior PDF on θ to average this derivative over all values of θ . The result is

$$\frac{d}{d\tilde{\theta}} \langle P_e(\theta, \tilde{\theta}) \rangle_{\theta} \Big|_{\tilde{\theta} = \theta \pm} = \mp \frac{1}{4} \left\langle \int_{-\infty}^{\infty} |\text{pr}'(\theta|\mathbf{g})| d\theta \right\rangle_{\mathbf{g}}. \quad (21)$$

The notation $\text{pr}'(\theta|\mathbf{g})$ indicates the derivative of $\text{pr}(\theta|\mathbf{g})$ with respect to θ . Thus, the integral on the right in this equation is the TV of the posterior PDF for θ . As we will see below, if the posterior PDF is unimodal, then the posterior TV is proportional to the height at the peak of the PDF. Higher peaks correspond to narrower posterior distributions, which, in turn, result in better estimates of the parameter of interest. This integral is then averaged over all data vectors to produce a figure of merit which we would expect to be associated with better

estimation performance overall. This equation shows that this average posterior TV is also connected with better performance for the task of detecting a small change in the parameter.

The function $P_e(\theta, \tilde{\theta})$ reaches its maximum value of 0.5 when $\tilde{\theta} = \theta$, and decreases as we move away from this point. Much like the function $\text{AUC}(\theta, \tilde{\theta})$ discussed above, $P_e(\theta, \tilde{\theta})$ is not differentiable at $\tilde{\theta} = \theta$, as a function of $\tilde{\theta}$, but the one-sided derivatives do exist on either side of this point. Thus, the average TV of the posterior PDF tells us how rapidly the MPE decreases as we move away from $\tilde{\theta} = \theta$. Another interpretation of this result is that the average TV of the posterior PDF gives us the average MPE for detecting a small change in the parameter θ .

The second result relates these same one-sided derivatives of the MPE to the Bayesian FI. Specifically we have

$$\left| \frac{d}{d\tilde{\theta}} \langle P_e(\theta, \tilde{\theta}) \rangle_{\theta} \Big|_{\tilde{\theta} = \theta \pm} \right| \leq \frac{1}{4} \sqrt{F}. \quad (22)$$

In the examples we will show that there are cases where we have equality between the two sides of this inequality, and in that sense the inequality is sharp. The next section shows the derivations of these two relations for those who are interested in seeing the details. For those who are not, this section can be skipped without any loss of understanding of the main results.

8. DERIVATION OF THE MAIN RESULTS

The ideal observer for the classification task in the Ziv-Zakai inequality can be formulated by defining a test statistic $t(\mathbf{g}|\theta, \tilde{\theta})$ by

$$t(\mathbf{g}|\theta, \tilde{\theta}) = \frac{\text{pr}(\tilde{\theta}|\mathbf{g})}{\text{pr}(\theta|\mathbf{g})} = \frac{\text{pr}(\mathbf{g}|\tilde{\theta}) \text{pr}(\tilde{\theta})}{\text{pr}(\mathbf{g}|\theta) \text{pr}(\theta)} = \frac{\text{pr}(\tilde{\theta})}{\text{pr}(\theta)} \Lambda(\mathbf{g}|\theta, \tilde{\theta}) \quad (23)$$

and declaring the hypothesis H_1 to be correct if $t(\mathbf{g}|\theta, \tilde{\theta}) > 1$. This is obviously equivalent to using the likelihood ratio test statistic and the threshold $y(\theta, \tilde{\theta})$ given above. Thus this classification scheme achieves the minimum possible value of the probability of error for this task. We will use the following notation for certain derivatives. For the test statistic we have

$$t'(\mathbf{g}|\theta) = \frac{d}{d\tilde{\theta}} t(\mathbf{g}|\theta, \tilde{\theta}) \Big|_{\tilde{\theta} = \theta}. \quad (24)$$

For the conditional and prior PDFs, we use

$$\text{pr}'(\mathbf{g}|\theta) = \frac{d}{d\tilde{\theta}} \text{pr}(\mathbf{g}|\tilde{\theta}) \Big|_{\tilde{\theta} = \theta} \quad (25)$$

and

$$\text{pr}'(\theta) = \frac{d}{d\theta} \text{pr}(\tilde{\theta}) \Big|_{\tilde{\theta} = \theta}. \quad (26)$$

These derivatives are related by the equation

$$t'(\mathbf{g}|\theta) = \frac{\text{pr}'(\mathbf{g}|\theta)}{\text{pr}(\mathbf{g}|\theta)} + \frac{\text{pr}'(\theta)}{\text{pr}(\theta)}. \quad (27)$$

This notation will make the derivation of the main results easier to follow.

The Bayesian FI can now be written as

$$F = \left\langle \left[t'(\mathbf{g}|\theta) \right]^2 \right\rangle_{\mathbf{g}|\theta} \Big|_{\theta}. \quad (28)$$

The function $t'(\mathbf{g}|\theta)$ can be viewed as a random variable t' since $\mathbf{g}$ is a random vector with conditional PDF $\text{pr}(\mathbf{g}|\theta)$, and θ is a random variable with PDF $\text{pr}(\theta)$. The mean of this random variable is given by

$$\left\langle t'(\mathbf{g}|\theta) \right\rangle_{\mathbf{g}|\theta} \Big|_{\theta} = \left\langle \left[\frac{\text{pr}'(\mathbf{g}|\theta)}{\text{pr}(\mathbf{g}|\theta)} \right] \right\rangle_{\mathbf{g}|\theta} \Big|_{\theta} + \frac{\text{pr}'(\theta)}{\text{pr}(\theta)} \Big|_{\theta} = \frac{\text{pr}'(\theta)}{\text{pr}(\theta)} \Big|_{\theta} = 0, \quad (29)$$

and we therefore we have $F = \text{var}(t')$.

For fixed $(\theta, \tilde{\theta})$ the function $t(\mathbf{g}|\theta, \tilde{\theta})$ can be viewed as a random variable t with conditional PDFs $\text{pr}_0(t|\theta, \tilde{\theta})$ and $\text{pr}_1(t|\theta, \tilde{\theta})$ under the two hypotheses H_0 and H_1 , respectively. In terms of these PDFs the MPE function $P_e(\theta, \tilde{\theta})$ can be written as

$$P_e(\theta, \tilde{\theta}) = \text{Pr}_0(\theta, \tilde{\theta}) \int_1^\infty \text{pr}_0(t|\theta, \tilde{\theta}) dt + \text{Pr}_1(\theta, \tilde{\theta}) \times \int_0^1 \text{pr}_1(t|\theta, \tilde{\theta}) dt. \quad (30)$$

We want to compute the derivative of this function with respect to $\tilde{\theta}$ evaluated at $\tilde{\theta} = \theta$. The magnitude of this derivative tells us how rapidly the MPE changes as $\tilde{\theta}$ moves away from θ . This in turn tells us how useful the data is for the Ziv-Zakai classification task when $\tilde{\theta}$ is close to θ .

Before we try to compute the derivative in question we need to explain the relation between the PDFs $\text{pr}_0(t|\theta, \tilde{\theta})$ and $\text{pr}_1(t|\theta, \tilde{\theta})$. The corresponding PDFs for the likelihood ratio satisfy the relation $\text{pr}_1(A|\theta, \tilde{\theta}) = A \text{pr}_0(A|\theta, \tilde{\theta})$. To see how this property translates to the PDFs $\text{pr}_0(t|\theta, \tilde{\theta})$ and $\text{pr}_1(t|\theta, \tilde{\theta})$ we consider two random variables w and x that are related by $w = cx$ for some constant c . In this case we have the standard relation $\text{pr}_w(w) = (1/c) \text{pr}_x(w/c)$. Now suppose we have a different PDF $\tilde{\text{pr}}_x(x)$ given by $\tilde{\text{pr}}_x(x) = x \text{pr}_x(x)$. The corresponding PDF for w is then given by

$$\tilde{\text{pr}}_w(w) = \frac{1}{c} \tilde{\text{pr}}_x\left(\frac{w}{c}\right) = \frac{w}{c^2} \text{pr}_x\left(\frac{w}{c}\right) = \frac{w}{c} \text{pr}_w(w). \quad (31)$$

Translating this result to the random variable t we have

$$\text{pr}_1(t|\theta, \tilde{\theta}) = \frac{\text{pr}(\theta)}{\text{pr}(\tilde{\theta})} t \text{pr}_0(t|\theta, \tilde{\theta}). \quad (32)$$

Using this result we may write the MPE function as

$$P_e(\theta, \tilde{\theta}) = \text{Pr}_0(\theta, \tilde{\theta}) \left[\int_1^\infty \text{pr}_0(t|\theta, \tilde{\theta}) dt + \int_0^1 t \text{pr}_0(t|\theta, \tilde{\theta}) dt \right]. \quad (33)$$

Using the normalization integral for $\text{pr}_0(t|\theta, \tilde{\theta})$ we can write

$$P_e(\theta, \tilde{\theta}) = \text{Pr}_0(\theta, \tilde{\theta}) \left[1 + \int_0^1 (t-1) \text{pr}_0(t|\theta, \tilde{\theta}) dt \right]. \quad (34)$$

Since the integral in this expression is an expectation, it can be written in terms of an expectation over the data vector $\mathbf{g}$ and we have

$$P_e(\theta, \tilde{\theta}) = \text{Pr}_0(\theta, \tilde{\theta}) \left[1 - \int_{\mathcal{D}} [1 - t(\mathbf{g}|\theta, \tilde{\theta})] \times \text{step}[1 - t(\mathbf{g}|\theta, \tilde{\theta})] \text{pr}(\mathbf{g}|\theta) d^M \mathbf{g} \right], \quad (35)$$

where $\mathcal{D}$ is the data domain. This is the form for the MPE of the Ziv-Zakai classification task that we will find most useful.

Now we suppose that $\theta > 0$. The case with negative θ will be similar. We will assume that the second derivative of $t(\mathbf{g}|\theta, \tilde{\theta})$ with respect to $\tilde{\theta}$ is continuous and note that $t(\mathbf{g}|\theta, \theta) = 1$. Using Taylor's theorem with remainder we find that, for θ small enough we can write

$$\begin{aligned} \text{step}[1 - t(\mathbf{g}|\theta, \theta + \Delta\theta)] &= \text{step}[-t'(\mathbf{g}|\theta) \Delta\theta] \\ &= \text{step}[-t'(\mathbf{g}|\theta)]. \end{aligned} \quad (36)$$

Therefore we have the expansion

$$P_e(\theta, \theta + \Delta\theta) = \left[\frac{1}{2} - \frac{\text{pr}'(\theta)}{4\text{pr}(\theta)} \Delta\theta \right] \times \left[1 + \int_{\mathcal{A}} [t'(\mathbf{g}|\theta) \Delta\theta] \text{pr}(\mathbf{g}|\theta) d^M \mathbf{g} \right] + \dots, \quad (37)$$

where $\mathcal{A}$ is the subset of data vectors in $\mathcal{D}$ satisfying $t'(\mathbf{g}|\theta) < 0$. Keeping only the zero and first-order terms we have

$$P_e(\theta, \theta + \Delta\theta) - \frac{1}{2} = \left[-\frac{\text{pr}'(\theta)}{4\text{pr}(\theta)} + \frac{1}{2} \int_{\mathcal{A}} t'(\mathbf{g}|\theta) \text{pr}(\mathbf{g}|\theta) d^M \mathbf{g} \right] \Delta\theta + \dots \quad (38)$$

We may write this with an integral over the whole data domain as

$$P_e(\theta, \theta + \Delta\theta) - \frac{1}{2} = \left[-\frac{\text{pr}'(\theta)}{4\text{pr}(\theta)} + \frac{1}{2} \int_{\mathcal{D}} \{1 - \text{step}[t'(\mathbf{g}|\theta)]\} \times t'(\mathbf{g}|\theta) \text{pr}(\mathbf{g}|\theta) d^M \mathbf{g} \right] \Delta\theta + \dots \quad (39)$$

Using the definition $t'(\mathbf{g}|\theta)$ this last expression simplifies to

$$P_e(\theta, \theta + \Delta\theta) - \frac{1}{2} = \left[\frac{\text{pr}'(\theta)}{4\text{pr}(\theta)} - \frac{1}{2} \int_{\mathcal{D}} \text{step}[t'(\mathbf{g}|\theta)] \times t'(\mathbf{g}|\theta) \text{pr}(\mathbf{g}|\theta) d^M \mathbf{g} \right] \Delta\theta + \dots \quad (40)$$

This will be the MPE expression we will use going forward.

Now we need to take an expectation over θ using the prior PDF $\text{pr}(\theta)$. The result is

$$\left\langle P_e(\theta, \theta + \Delta\theta) - \frac{1}{2} \right\rangle_{\theta} = -\frac{1}{2} \left\langle \int_{\mathcal{D}} \text{step}[t'(\mathbf{g}|\theta)] t'(\mathbf{g}|\theta) \text{pr}(\mathbf{g}|\theta) d^M \mathbf{g} \right\rangle_{\theta} \Delta\theta + \dots \quad (41)$$

In terms of expectations we can now write

$$\langle 1 - 2P_e(\theta, \theta + \Delta\theta) \rangle_{\theta} = \langle \langle t' \text{step}(t') \rangle_{\mathbf{g}|\theta} \rangle_{\theta} \Delta\theta + \dots \quad (42)$$

By considering the $\theta < 0$ case we can summarize the results in terms of one-side derivatives as

$$\left\langle \frac{d}{d\tilde{\theta}} [1 - 2P_e(\theta, \tilde{\theta})] \right\rangle_{\tilde{\theta} = \theta \pm} \Big|_{\theta} = \pm \langle \langle t' \text{step}(t') t \rangle_{\mathbf{g}|\theta} \rangle_{\theta}. \quad (43)$$

Since the mean of t' is zero we have the final result

$$\left\langle \frac{d}{d\tilde{\theta}} [1 - 2P_e(\theta, \tilde{\theta})] \right\rangle_{\tilde{\theta} = \theta \pm} \Big|_{\theta} = \pm \frac{1}{2} \langle \langle |t'| \rangle_{\mathbf{g}|\theta} \rangle_{\theta}. \quad (44)$$

This is similar to the result that we discussed in Section 3 relating FI to the ideal-observer AUC for this classification task. For fixed θ the function $1 - 2P_e(\theta, \tilde{\theta})$ reaches a minimum value of zero at $\tilde{\theta} = \theta$ but it is not differentiable there. The slope as we move away from this singularity is determined by the mean value of the random variable $|t'|$. Now, using the Schwarz inequality we can bring in the Bayesian FI as follows:

$$\left| \left\langle \frac{d}{d\tilde{\theta}} [1 - 2P_e(\theta, \tilde{\theta})] \right\rangle_{\tilde{\theta} = \theta \pm} \Big|_{\theta} \right| \leq \frac{1}{2} \sqrt{F}. \quad (45)$$

We will see in the examples below that it is possible to have equality in this relation. In terms of the MPE for the Ziv-Zakai task directly, we may write the one-sided derivatives as

$$\frac{d}{d\tilde{\theta}} \langle P_e(\theta, \tilde{\theta}) \rangle_{\theta} \Big|_{\tilde{\theta} = \theta_{\pm}} = \mp \frac{1}{4} \langle \langle |t'| \rangle_{\mathbf{g}} | \theta \rangle_{\theta}. \quad (46)$$

The Bayesian FI inequality then has the form

$$\left| \frac{d}{d\tilde{\theta}} \langle P_e(\theta, \tilde{\theta}) \rangle_{\theta} \Big|_{\tilde{\theta} = \theta_{\pm}} \right| \leq \frac{1}{4} \sqrt{F}. \quad (47)$$

This inequality completes the circle relating MPE for the Ziv-Zakai classification task, the EMSE for a parameter estimator, and the Bayesian FI. It also provides a new interpretation of the Bayesian FI in terms of the average MPE for the task of detecting a small change in a parameter.

As a final note in this section we can rearrange the expectations in the mean of $|t'|$ and find that

$$\langle \langle |t'| \rangle_{\mathbf{g}} | \theta \rangle_{\theta} = \langle \langle |t'| \rangle_{\theta} | \mathbf{g} \rangle_{\mathbf{g}} = \left\langle \int_{-\infty}^{\infty} |\text{pr}'(\theta | \mathbf{g})| d\theta \right\rangle_{\mathbf{g}}. \quad (48)$$

This last expectation is the average value of the TV of the posterior PDF $\text{pr}(\theta | \mathbf{g})$. It is this quantity which governs the behavior of $\langle P_e(\theta, \tilde{\theta}) \rangle_{\theta}$ when $\tilde{\theta}$ is close to θ . From standard results about the total variation we then have

$$\frac{d}{d\tilde{\theta}} \langle P_e(\theta, \tilde{\theta}) \rangle_{\theta} \Big|_{\tilde{\theta} = \theta_{\pm}} \leq -\frac{1}{4} \left\langle \frac{1}{N} \sum_{n=1}^N |\text{pr}(\theta_n | \mathbf{g}) - \text{pr}(\theta_{n-1} | \mathbf{g})| \right\rangle_{\mathbf{g}} \quad (49)$$

and

$$\frac{d}{d\tilde{\theta}} \langle P_e(\theta, \tilde{\theta}) \rangle_{\theta} \Big|_{\tilde{\theta} = \theta_{-}} \geq \frac{1}{4} \left\langle \frac{1}{N} \sum_{n=1}^N |\text{pr}(\theta_n | \mathbf{g}) - \text{pr}(\theta_{n-1} | \mathbf{g})| \right\rangle_{\mathbf{g}}. \quad (50)$$

For large values of N the sums can give us a good approximation to the TV of the posterior PDF when it cannot be computed analytically. The TV of the posterior PDF can be thought of as an average figure of merit for the task of detecting a small change in a parameter that is also related, via the Bayesian FI and the van Trees inequality, to the EMSE of an estimator for the same parameter.

9. VECTOR VERSION FOR BAYESIAN FIM AND MPE

The results of the previous section can also be extended to p -dimensional vector parameters. The Bayesian FIM is the $p \times p$ matrix

$$\mathbf{F} = \left\langle \left[\nabla_{\theta} \ln \text{pr}(\theta | \mathbf{g}) \right] \left[\nabla_{\theta} \ln \text{pr}(\theta | \mathbf{g}) \right]^{\dagger} \right\rangle_{\mathbf{g} | \theta} \Big|_{\theta}. \quad (51)$$

If $\mathbf{u}$ is a unit vector in the parameter space, then we can show that the one-sided directional derivatives satisfy

$$\frac{d}{dt} \langle P_e(\theta, \theta + t\mathbf{u}) \rangle_{\theta} \Big|_{t=0+} \geq -\frac{1}{4} \sqrt{\mathbf{u}^{\dagger} \mathbf{F} \mathbf{u}} \quad (52)$$

and

$$\frac{d}{dt} \langle P_e(\theta, \theta + t\mathbf{u}) \rangle_{\theta} \Big|_{t=0-} \leq \frac{1}{4} \sqrt{\mathbf{u}^{\dagger} \mathbf{F} \mathbf{u}}. \quad (53)$$

Therefore, the components of the Bayesian FIM are related to the change in the average MPE $\langle P_e(\theta, \theta') \rangle_{\theta}$ as θ' moves away from θ . There are vector versions of the van Trees inequality and the Ziv-Zakai inequality, so in the vector case also the EMSE, the Bayesian FIM, and the MPE for the task of detecting a change in the parameter vector are all related.

10. EXAMPLES AND IMAGING APPLICATIONS

Since we are really only interested in the magnitudes of the one-sided derivatives of the MPE, we will summarize the results derived in Section 8 as follows:

$$\left| \frac{d}{d\tilde{\theta}} \langle P_e(\theta, \tilde{\theta}) \rangle_{\theta} \Big|_{\tilde{\theta} = \theta \pm} \right| = \frac{1}{4} \left\langle \int_{-\infty}^{\infty} |\text{pr}'(\theta | \mathbf{g})| d\theta \right\rangle_{\mathbf{g}} \leq \frac{1}{4} \sqrt{F}. \quad (54)$$

Suppose that the posterior PDF $\text{pr}(\theta | \mathbf{g})$ is supported on the (possibly infinite) interval between a and b . If this PDF is unimodal with mode x , then the TV of the posterior can be written as

$$\int_a^b |\text{pr}'(\theta | \mathbf{g})| d\theta = \int_a^x \text{pr}'(\theta | \mathbf{g}) d\theta - \int_x^b \text{pr}'(\theta | \mathbf{g}) d\theta. \quad (55)$$

By the fundamental theorem of calculus we then have

$$\int_a^b |\text{pr}'(\theta | \mathbf{g})| d\theta = 2\text{pr}(x | \mathbf{g}) + \text{pr}(b | \mathbf{g}) - \text{pr}(a | \mathbf{g}). \quad (56)$$

Thus, the TV for the posterior PDF is easy to calculate in this case if we have an analytic expression for this PDF.

We consider two examples of possible unimodal posterior PDFs. The first is a normal distribution:

$$\text{pr}(\theta|m, \sigma) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp \left[-\frac{1}{2\sigma^2}(\theta - m)^2 \right]. \quad (57)$$

Since this PDF is unimodal with a maximum at $\theta = m$, we have immediately that the posterior TV is given by

$$\int_{-\infty}^{\infty} |\text{pr}'(\theta|m, \sigma)| d\theta = \frac{1}{\sigma} \sqrt{\frac{2}{\pi}}. \quad (58)$$

The corresponding value for the Bayesian FI is given by

$$\left\langle \left[\frac{\text{pr}'(\theta|m, \sigma)}{\text{pr}(\theta|m, \sigma)} \right]^2 \right\rangle_{\theta|m, \sigma} = \frac{1}{\sigma^2}. \quad (59)$$

Thus, the ratio of the posterior TV to the square root of the Bayesian FI is $\sqrt{2/\pi}$. This number is less than unity as expected.

The second example of a possible posterior PDF is the exponential distribution:

$$\text{pr}(\theta|\beta) = \beta \exp(-\beta\theta). \quad (60)$$

In this case the posterior TV is given by

$$\int_0^{\infty} |\text{pr}'(\theta|\beta)| d\theta = \beta, \quad (61)$$

while the Bayesian FI is

$$\left\langle \left[\frac{\text{pr}'(\theta|\beta)}{\text{pr}(\theta|\beta)} \right]^2 \right\rangle_{\theta|\alpha, \beta} = \beta^2. \quad (62)$$

The ratio of the posterior TV to the square root of the Bayesian FI for this example is unity. Since there are many cases where an exponential distribution is a posterior PDF, this example shows that it is possible to have equality in the inequality in Eq. (54). Therefore we may also have equality between one fourth of the square root of the Bayesian FI and the magnitude of the one-sided derivatives of the MPE for the task of detecting a small change in a parameter.

Now we consider two imaging applications where we have a conditional distribution $\text{pr}(\mathbf{g}|\boldsymbol{\theta})$ describing the statistical properties of the imaging system noise and a prior distribution $\text{pr}(\boldsymbol{\theta})$ on the parameter vector of interest. For the first imaging application we have a scalar parameter and Gaussian statistics for the conditional PDF $\text{pr}(\mathbf{g}|\boldsymbol{\theta})$ and the prior PDF $\text{pr}(\boldsymbol{\theta})$. Specifically, we fix a unit vector $\mathbf{s}$ in data space and assume that the conditional PDF is given by

$$\text{pr}(\mathbf{g}|\theta) = \frac{1}{\sqrt{2\pi \det \mathbf{K}}} \exp \left[-\frac{1}{2}(\mathbf{g} - \theta \mathbf{s})^\dagger \mathbf{K}^{-1}(\mathbf{g} - \theta \mathbf{s}) \right]. \quad (63)$$

This PDF describes an imaging situation where $\mathbf{g} = \theta \mathbf{s} + \mathbf{n}$, the signal shape $\mathbf{s}$ is known, and $\mathbf{n}$ is a background of textured Gaussian noise. In the estimation context we are trying to estimate the amplitude θ of the signal. In the detection context we are trying to detect a change in this amplitude. We may decompose the data vector as $\mathbf{g} = \mathbf{g}_\perp + \mathbf{g}_\parallel$ with $\mathbf{g}_\perp^\dagger \mathbf{K}^{-1} \mathbf{s} = 0$ and $\mathbf{g}_\parallel = g_\parallel \mathbf{s}$. Then the conditional PDF may be written as

$$\text{pr}(\mathbf{g}|\theta) = \frac{1}{\sqrt{2\pi \det \mathbf{K}}} \exp \left[-\frac{1}{2}(g_\parallel - \theta)^\dagger \mathbf{s}^\dagger \mathbf{K}^{-1} \mathbf{s} (g_\parallel - \theta) \right] \times \exp \left[-\frac{1}{2} \mathbf{g}_\perp^\dagger \mathbf{K}^{-1} \mathbf{g}_\perp \right]. \quad (64)$$

We will assume a Gaussian prior PDF:

$$\text{pr}(\theta) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp \left[-\frac{1}{2\sigma^2}(\theta - \mu)^2 \right]. \quad (65)$$

The posterior PDF is also a Gaussian and all that we need is its variance σ_p^2 , which is given by

$$\frac{1}{\sigma_p^2} = \mathbf{s}^\dagger \mathbf{K}^{-1} \mathbf{s} + \frac{1}{\sigma^2}. \quad (66)$$

Now for the posterior TV we have

$$\int_{-\infty}^{\infty} |\text{pr}'(\theta|\mathbf{g})| d\theta = \sqrt{\frac{2}{\pi\sigma_p^2}} = \sqrt{\frac{2}{\pi} \left(\mathbf{s}^\dagger \mathbf{K}^{-1} \mathbf{s} + \frac{1}{\sigma^2} \right)}. \quad (67)$$

Since this quantity does not depend on $\mathbf{g}$, we have the average posterior TV:

$$\langle \langle |\text{pr}'(\theta|\mathbf{g})| \rangle_{\mathbf{g}|\theta} \rangle_\theta = \sqrt{\frac{2}{\pi} \left(\mathbf{s}^\dagger \mathbf{K}^{-1} \mathbf{s} + \frac{1}{\sigma^2} \right)}. \quad (68)$$

The Bayesian FI is given by

$$F = \mathbf{s}^\dagger \mathbf{K}^{-1} \mathbf{s} + \frac{1}{\sigma^2}. \quad (69)$$

Therefore, the ratio of the average posterior TV and the square root of the Bayesian FI is $\sqrt{2/\pi}$.

In this application we consider a conditional PDF of the following form:

$$\text{pr}(\mathbf{g}|\boldsymbol{\theta}) = \frac{1}{\sqrt{2\pi \det \mathbf{K}_n}} \exp \left[-\frac{1}{2}(\mathbf{g} - \mathbf{H}\boldsymbol{\theta})^\dagger \mathbf{K}_n^{-1} (\mathbf{g} - \mathbf{H}\boldsymbol{\theta}) \right]. \quad (70)$$

We may think of the $M \times N$ matrix $\mathbf{H}$ as representing an imaging system acting on the N -dimensional object vector $\boldsymbol{\theta}$ and generating the M -dimensional noisy data vector $\mathbf{g}$, where the noise is described by correlated Gaussian statistics. The vector $\boldsymbol{\theta}$ may, for example, be a voxelized representation of an object in a tomographic imaging system, and the matrix $\mathbf{H}$ a model of that system derived from calibration, an analytic model, or a Monte Carlo model. The estimation task in this case is a voxelized object reconstruction of the vector $\boldsymbol{\theta}$. The related detection task is to detect a change in this parameter vector. We will assume that $M < N$ and that the matrix $\mathbf{H}$ is full rank. If $\mathbf{H}^+$ is the pseudoinverse of $\mathbf{H}$, then these assumptions imply that $\mathbf{H}\mathbf{H}^+ = \mathbf{I}$. We may therefore write the conditional PDF as

$$\text{pr}(\mathbf{g}|\boldsymbol{\theta}) = \frac{1}{\sqrt{2\pi \det \mathbf{K}_n}} \times \exp \left[-\frac{1}{2}(\boldsymbol{\theta} - \mathbf{H}^+ \mathbf{g})^\dagger \mathbf{H}^\dagger \mathbf{K}_n^{-1} \mathbf{H} (\boldsymbol{\theta} - \mathbf{H}^+ \mathbf{g}) \right]. \quad (71)$$

We assume the prior PDF is also Gaussian and given by

$$\text{pr}(\boldsymbol{\theta}) = \frac{1}{\sqrt{2\pi \det \mathbf{K}}} \exp \left[-\frac{1}{2}(\boldsymbol{\theta} - \boldsymbol{\mu})^\dagger \mathbf{K}^{-1} (\boldsymbol{\theta} - \boldsymbol{\mu}) \right]. \quad (72)$$

It is now easy to see that the posterior PDF is also Gaussian with a covariance matrix given by $\mathbf{K}_p^{-1} = \mathbf{K}_\theta^{-1} + (\mathbf{H}^\dagger \mathbf{K}_n^{-1} \mathbf{H})$. If we define a directional derivative for the unit vector $\mathbf{u}$ in parameter space by

$$D_{\mathbf{u}} \text{pr}(\boldsymbol{\theta}|\mathbf{g}) = \frac{d}{dt} \text{pr}(\boldsymbol{\theta} + t\mathbf{u}|\mathbf{g})|_{t=0}, \quad (73)$$

then the vector analog of the posterior TV for the case of a scalar parameter is given by

$$\int_{\mathbb{R}^N} |D_{\mathbf{u}} \text{pr}(\boldsymbol{\theta}|\mathbf{g})| d^N \boldsymbol{\theta} = \int_{\mathbb{R}^N} |\mathbf{u}^\dagger \mathbf{K}_p^{-1} (\boldsymbol{\theta} - \boldsymbol{\mu}_p)| \text{pr}(\boldsymbol{\theta}|\mathbf{g}) d^N \boldsymbol{\theta}. \quad (74)$$

This integral can be computed and reduces to

$$\int_{\mathbb{R}^N} |D_{\mathbf{u}} \text{pr}(\boldsymbol{\theta}|\mathbf{g})| d^N \boldsymbol{\theta} = \sqrt{\frac{2}{\pi} \mathbf{u}^\dagger \mathbf{K}_p^{-1} \mathbf{u}}. \quad (75)$$

As in the scalar case, this number is the magnitude of the one-sided directional derivative of $P_{\mathcal{A}}(\boldsymbol{\theta}, \boldsymbol{\theta}')$ when $\boldsymbol{\theta}'$ is moving away from $\boldsymbol{\theta}$ in the direction $\mathbf{u}$. The square root of the $\mathbf{u}$ component of the Bayesian FIM is given by

$$\sqrt{\mathbf{u}^\dagger \mathbf{F} \mathbf{u}} = \sqrt{\mathbf{u}^\dagger \mathbf{K}_p^{-1} \mathbf{u}}. \quad (76)$$

The ratio of these two quantities is once again $\sqrt{2/\pi}$. This reflects the fact that the multivariate version of the posterior TV in the direction $\mathbf{u}$ will always be less than or equal to $\sqrt{\mathbf{u}^\dagger \mathbf{F} \mathbf{u}}$, a fact which can be proved using the same methods used above for the scalar case.

11. CONCLUSION

The relation between the FI and the ideal-observer AUC described above relates the FI to the ability of the ideal observer to detect a small change in a scalar parameter that is affecting the statistics of the data vector. The relation between the FIM and the ideal-observer AUC is similar except that we are trying to detect a change in a vector parameter. In both cases the AUC is approximately proportional to the square root of the relevant component of the FIM for small changes in the parameters. In this work we extended these results to the Bayesian FI for scalar parameters and the Bayesian FIM for vector parameters. The ideal-observer AUC and the FIM do not depend on the prior probabilities for the signal present and signal absent hypotheses. The new element in the Bayesian approach is a prior on the parameters governing the statistics of the data, which can be used to define these prior probabilities.

This extension is based on a task introduced in the Ziv-Zakai inequality where the ideal observer is trying to detect a change in a parameter and the prior probabilities for the two hypotheses are determined by the prior PDF on the parameter. In this case the AUC for the ideal observer is no longer relevant, and it is the probability of error for the ideal observer, the MPE, that becomes the detection figure of merit. The Ziv-Zakai inequality relates the MPE to the EMSE for any estimator of the parameter. In this work we related the average MPE for small deviations in the parameter to the Bayesian FIM via an inequality. An example shows that this inequality will be equality for certain posterior PDFs, so in this sense the inequality is sharp. An intermediate quantity in the derivation of this result is the average posterior TV, which is related to the small-deviation MPE as the first term in a Taylor series expansion. This relation is similar to the relation between the ideal-observer AUC and the FI discussed above.

The results discussed in this work further elucidate the connections between estimating a parameter and detecting a change in that parameter. An imaging system optimized for one of these tasks will probably be optimized for the other. In particular, if we are using the FI or the Bayesian FI for optimization on an estimation task, then we are also optimizing for the task of detecting a small change in the parameter of interest. The results in this paper and others [4,5,6,9,12,13] also show that the well-known measures of information, ideal-observer AUC, MPE, FI, Bayesian FI, and Shannon information, are all related to each other in ways that are not always obvious. We may also now add the posterior TV to this list as a measure of information related to both detection and estimation tasks.

Funding.

National Institute of Biomedical Imaging and Bioengineering (P41-EB002035, R01-EB000803).

REFERENCES

1. Shao J, Mathematical Statistics (Springer, 1999).
2. Barrett HH and Myers KJ, Foundations of Image Science (Wiley, 2004).
3. Barrett H, Abbey C, and Clarkson E, "Objective assessment of image quality. III. ROC metrics, ideal observers, and likelihood-generating functions," J. Opt. Soc. Am. A 15, 1520–1535 (1998).
4. Clarkson E and Shen F, "Fisher information and surrogate figures of merit for the task-based assessment of image quality," J. Opt. Soc. Am. A 27, 2313–2326 (2010).
5. Shen F and Clarkson E, "Using Fisher information to approximate ideal observer performance on detection tasks for lumpy-background images," J. Opt. Soc. Am. A 23, 2406–2414 (2006).
6. Clarkson E, "Asymptotic ideal observers and surrogate figures of merit for signal detection with list-mode data," J. Opt. Soc. Am. A 29, 2204–2216 (2012).
7. Ziv J and Zakai M, "Some lower bounds on signal parameter," IEEE Trans. Inf. Theory 15, 386–391 (1969).
8. Bell K, Steinberg Y, Ephraim Y, and Van Trees H, "Extended Ziv-Zakai lower bound for vector parameter estimation," IEEE Trans. Inf. Theory 43, 624–637 (1997).
9. Clarkson E, "The relation between Bayesian Fisher information and Shannon information for detecting a change in a parameter," J. Opt. Soc. Am. A 36, 1209–1214 (2019).
10. van Trees HL, Detection, Estimation and Modulation Theory, Part 1 (Wiley, 1968).
11. Gill RD and Levit BY, "Applications of the van Trees inequality: a Bayesian Cramér–Rao bound," Bernoulli 1, 59–79 (1995).
12. Clarkson E and Cushing J, "Shannon information and ROC analysis in imaging," J. Opt. Soc. Am. A 32, 1288–1301 (2015).
13. Clarkson E and Cushing J, "Shannon information and receiver operating characteristic analysis for multi-class classification in imaging," J. Opt. Soc. Am. A 33, 930–937 (2016).