

# Behavior of the decision variables of the three-class ideal observer for univariate trinormal data

Darrin C. Edwards\* and Charles E. Metz

Department of Radiology, The University of Chicago, Chicago, IL 60637

## ABSTRACT

We are attempting to extend receiver operating characteristic (ROC) analysis to tasks with more than two classes. This is difficult because of the rapid increase in complexity, in general, of both observer behavior and evaluation of its performance, as the number of classes involved increases. Many researchers have proposed addressing this complexity by imposing simplifications on the model; for example, by using univariate data rather than the bivariate data employed by the general three-class ideal observer. We have investigated a univariate trinormal model for the underlying data of a three-class ideal observer. Although a reasonably complete description of the ideal observer's behavior in this case is attainable, this behavior is more complicated than might intuitively be expected.

**Keywords:** ROC methodology, observer performance evaluation, ideal observers

## 1. INTRODUCTION

We are attempting to extend receiver operating characteristic (ROC) analysis to tasks with more than two classes. This is difficult because of the rapid increase in complexity, in general, of both observer behavior and evaluation of its performance, as the number of classes involved increases. In particular, we have shown that the hypervolume under the ROC hypersurface, an obvious generalization of the area under the ROC curve (AUC), does not in fact yield a useful performance metric in tasks with three or more classes.<sup>1</sup> To address this problem, we have proposed a novel performance metric, called the “surface-averaged expected cost” (SAEC), motivated by straightforward considerations of the expected utility (or, equivalently, the “costs”) of decisions made by an observer.<sup>2-4</sup>

Although we have theoretical grounds for hoping that this proposed performance metric will not suffer from the same difficulties as the naive hypervolume performance metric, it nevertheless shares one assumption in common: that the ROC hypersurface in, for example, a three-class classification task is five-dimensional over some non-zero region of support in the domain of the six-dimensional ROC space in which that hypersurface is defined. Although this assumption might seem reasonable given that the ideal observer in a three-class task makes decisions based on a set of five independent decision criteria,<sup>5</sup> it cannot be denied that intuition has often proved an unreliable guide in generalizing from the two-class classification task to tasks with three or more classes. He and Frey, in particular, have conjectured that even in the case of the three-class ideal observer, the ROC hypersurface might not have five degrees of freedom anywhere.<sup>6</sup> If this conjecture were proved true, our proposed performance metric would be just as useless as the naive hypervolume metric: it would always be zero for any observer. Addressing this conjecture would require calculation of operating points on the ideal observer's ROC hypersurface in sufficient quantity and detail to allow thorough analysis of the properties of that hypersurface—in particular, properties of the set of vectors tangent to the hypersurface at any given operating point. Unfortunately, this problem has so far proved intractable given the forms of the integrals involved;<sup>5</sup> as a first step in this direction, however, one might reasonably consider imposing constraints or other simplifying assumptions on the behavior of the ideal observer<sup>7</sup> or on the underlying data itself.

One such simplifying assumption would be the consideration of univariate underlying data; recall that for multivariate underlying data, the three-class ideal observer makes use of bivariate decision variable data, namely a pair of likelihood ratios.<sup>5</sup> Note that other researchers have investigated univariate three-class observers, including

---

\*Correspondence: E-mail: d-edwards@uchicago.edu; Telephone: 773 834 5094; Fax: 773 702 0371

non-ideal observers. For example, work by Scurfield<sup>8</sup> assumed not only that the underlying observational data were univariate, but also that the observer’s decision rule took on a particular form, defined by a pair of thresholds on the underlying data; observations less than the lower threshold are assigned to one class, those between the two thresholds to a second class, and those above the higher threshold to a third class. This is a straightforward generalization of two-class classification models (such as the conventional<sup>9</sup> and proper<sup>10</sup> binormal models; recall that the latter is motivated by an ideal observer approach to the underlying data model of the former) in which a single threshold is used to separate univariate data into two classes. While more complicated observer behavior is certainly conceivable, it is perhaps surprising that the three-class ideal observer itself is not limited to such a decision strategy. For normally distributed underlying data, in particular, we have been able to show<sup>11,12</sup> that the ideal observer can, depending on the values of its five decision criteria<sup>5</sup> and the distributional parameters of the underlying data, operate by dividing the underlying data space into as many as seven regions (*i.e.*, by partitioning the underlying data space with as many as six thresholds).

In the present work, rather than repeat the detailed derivation of our results, which we have already provided elsewhere,<sup>11,12</sup> we review the results briefly while attempting to provide a more intuitive motivation for their counterintuitive aspects by comparing them with the behavior of the ideal observer for two-class normal (so-called “binormal”) underlying data. We go on to illustrate, in the three-class task for univariate trinormal underlying data, the range of different possible decision strategies (boundaries in LR space and corresponding thresholds in underlying data space) that can be obtained even for a fixed underlying data model (a given set of means and variances of the normal underlying data distributions). We then review briefly some of the implications of these results for the dimensionality properties of the ROC hypersurface of the three-class ideal observer for univariate trinormal underlying data. Our work demonstrates that although a reasonably complete description of the ideal observer’s behavior can be obtained in this special case, aspects of the ideal observer’s behavior are more complicated than might intuitively be expected.

## 2. THE IDEAL OBSERVER FOR UNIVARIATE BINORMAL DATA (A REVIEW)

The two-class classification task is generally modeled as an underlying data variable  $\mathbf{x}$  which is drawn from two probability density functions (PDFs)  $p(x|\pi_1)$  (the “actually positive” class) and  $p(x|\pi_2)$  (the “actually negative” class). Under very general assumptions, the decision strategy itself is modeled as comparing a decision variable, some function  $\mathbf{y} = g(\mathbf{x})$ , with a threshold  $y_0$ ; if the decision variable for a given observation is greater than this threshold, then the observation is assigned to class  $\pi_1$ , otherwise it is assigned to class  $\pi_2$ .

The simplest non-trivial decision variable is the data variable  $\mathbf{x}$  itself. One example of such a model is the well-known conventional binormal model,<sup>9</sup> for which

$$p(x|\pi_1) = N(x; \mu_1 = a/b, \sigma_1^2 = 1/b) \tag{1}$$

and

$$p(x|\pi_2) = N(x; \mu_2 = 0, \sigma_2^2 = 1) \tag{2}$$

are the distributions of the underlying data, and  $a$  and  $b$  are a pair of model parameters related to the relative mean and relative width of the distributions. An example of this model with  $a = 1$  and  $b = 1$ , and with a particular decision threshold of  $x_0 = 1/2$ , is shown in Fig. 1.

Of course, it is also well known that optimal performance in such a two-class classification task is achieved when the decision variable used is the likelihood ratio, defined as

$$\text{LR}(x) \equiv \frac{p(x|\pi_1)}{p(x|\pi_2)}. \tag{3}$$

For the binormal model, this can be explicitly evaluated using Eqs. 1 and 2 to obtain

$$\text{LR}(x) = \frac{1}{\sigma_1} \exp \left\{ -\frac{1}{2} \left[ \left( \frac{x - \mu_1}{\sigma_1} \right)^2 - x^2 \right] \right\}; \tag{4}$$




Figure 1. An example of the conventional binormal model with model parameters  $a = 1$ ,  $b = 1$ , and decision threshold  $x_0 = 1/2$ .

the ideal observer compares this function of the underlying data with a threshold  $LR_0$ , and it is this decision strategy which is assumed under the “proper” binormal model.<sup>10</sup> For the special case  $\sigma_1 = 1$ , it can be shown that a given threshold  $LR_0$  corresponds to a single threshold  $x_0$ , yielding a decision rule equivalent to that shown in Fig. 1. It is also fairly straightforward to show<sup>10</sup> that, for  $\sigma_1 < 1$ , deciding an observation belongs to class  $\pi_1$  if  $LR > LR_0$  is equivalent to deciding it belongs to class  $\pi_1$  if

$$x_1 < \mathbf{x} < x_2, \tag{5}$$

where  $x_1$  and  $x_2$  are the solutions of the quadratic equation

$$\left(\frac{x - \mu_1}{\sigma_1}\right)^2 - x^2 = -2 \log(\sigma_1 LR_0). \tag{6}$$

(For  $\sigma_1 > 1$ , the sense of the test is reversed; *i.e.*, class  $\pi_1$  is decided if  $\mathbf{x} < x_1$  or  $x_2 < \mathbf{x}$ . Because we are free, in the case of unequal variances, to label the classes such that class  $\pi_1$  has the smaller variance, this distinction is of no importance.) An example of this model with  $\mu_1 = 1$ ,  $\sigma_1 = 1/2$ , and a decision threshold of  $LR_0 = 1$ , is shown in Fig. 2. The resulting ROC operating point coordinates would be

$$TPF(LR_0) = \int_{x_1}^{x_2} p(x|\pi_1) dx \tag{7}$$

and

$$FPF(LR_0) = \int_{x_1}^{x_2} p(x|\pi_2) dx, \tag{8}$$

where TPF is the true-positive fraction and FPF is the false-positive fraction. (In a three-class classification task, it is convenient to take the set of six misclassification probabilities to be the coordinates of the ROC space.)

We thus see that even in the two-class case, the representation of a decision rule in the underlying data space can be more complicated than the simple  $\mathbf{y} = g(\mathbf{x}) > y_0$  decision rule. This observation has practical consequences, because which representation (underlying data space or decision variable space) is the more “natural” may depend on the application of interest. (Indeed, in the development of the PROPROC algorithm implementing the proper binormal model,<sup>10</sup> when calculating ROC coordinates, it was found more sensible computationally to first calculate the thresholds  $x_1$  and  $x_2$  above and integrate the corresponding normal PDFs, rather than first obtaining the actual PDFs of  $LR$  from Eq. 4<sup>13</sup> and then integrating these densities against the single threshold  $LR_0$ .)




Figure 2. An example of the proper binormal model with model parameters  $\mu_1 = 1$ ,  $\sigma_1 = 1/2$ , and decision threshold  $\text{LR}_0 = 1$ .

### 3. THE IDEAL OBSERVER FOR UNIVARIATE TRINORMAL DATA

In a three-class classification task with univariate underlying data, it is very natural to consider a straightforward generalization of the models of the preceding section. That is, we now have an underlying data variable  $\mathbf{x}$  drawn from three PDFs:  $p(x|\pi_1)$ ,  $p(x|\pi_2)$ , and  $p(x|\pi_3)$ . For concreteness, we will again assume these to be normal distributions, normalizing so that class  $\pi_3$  is a standard normal distribution:

$$p(x|\pi_1) = N(x; \mu_1, \sigma_1^2) \quad (9)$$

$$p(x|\pi_2) = N(x; \mu_2, \sigma_2^2) \quad (10)$$

$$p(x|\pi_3) = N(x; \mu_3 = 0, \sigma_3^2 = 1). \quad (11)$$

Again, it is natural to consider, at least at first, an obvious generalization of the decision strategy depicted in Fig. 1. That is, since we now have three instead of two classes, we take our decision rule to be a simple pair of thresholds  $x_1$  and  $x_2$ , with the first separating class  $\pi_2$  and  $\pi_3$ , and the second separating class  $\pi_3$  and  $\pi_1$ . An example of such a decision strategy, equivalent to that investigated by Scurfield,<sup>8</sup> is shown in Fig. 3 with  $\mu_1 = 1$ ,  $\sigma_1^2 = 1$ ,  $\mu_2 = -1$ ,  $\sigma_2^2 = 1$ , and with decision thresholds  $x_1 = -1/2$  and  $x_2 = 1/2$ .

The ideal observer in such a univariate task makes decisions based on a pair of likelihood ratio decision variables given by

$$\text{LR}_1 \equiv \frac{p(x|\pi_1)}{p(x|\pi_3)} \quad (12)$$

$$\text{LR}_2 \equiv \frac{p(x|\pi_2)}{p(x|\pi_3)}, \quad (13)$$

which, for the trinormal underlying data model of Eqs. 9–11, yield

$$\text{LR}_1 = \frac{1}{\sigma_1} \exp \left\{ -\frac{1}{2} \left[ \left( \frac{x - \mu_1}{\sigma_1} \right)^2 - x^2 \right] \right\} \quad (14)$$

$$\text{LR}_2 = \frac{1}{\sigma_2} \exp \left\{ -\frac{1}{2} \left[ \left( \frac{x - \mu_2}{\sigma_2} \right)^2 - x^2 \right] \right\}. \quad (15)$$

Since the data are univariate, this pair of equations actually gives  $\text{LR}_2$  as a relation of  $\text{LR}_1$ , which we will call the likelihood ratio curve. The ideal observer's decision strategy is based on a set of three decision boundary




Figure 3. An example of a naive generalization of the conventional binormal model to three classes, with model parameters  $\mu_1 = 1, \sigma_1^2 = 1, \mu_2 = -1, \sigma_2^2 = 1$ , and with decision thresholds  $x_1 = -1/2$  and  $x_2 = 1/2$ .




Figure 4. The likelihood ratio curve when the underlying data are drawn from unit-variance normal distributions with parameters  $\mu_1 = 1, \mu_2 = -1$ , and  $\mu_3 = 0$ , and where the decision boundary line segments are shown with dashed lines. An expression such as “ $\pi_1$ ” indicates a region where a particular decision is made.

line segments placed in the  $\vec{LR}$  plane, subject to various restrictions on the slopes and axis intercepts of those line segments.<sup>5,14</sup> For the specific underlying data model illustrated in Fig. 3, one can show that

$$LR_2 = \exp\left(-\frac{1}{2}\right) \left[LR_1 \exp\left(\frac{1}{2}\right)\right]^{-1}, \quad (16)$$

a power law relation between  $LR_2$  and  $LR_1$  with negative exponent. This is illustrated in Fig. 4 for a specific choice of decision criteria such that the decision boundaries are  $LR_2 = LR_1$ ,  $LR_1 = 1$ , and  $LR_2 = 1$ , which we refer to as the “1-vs.-2”, “1-vs.-3”, and “2-vs.-3” boundaries, respectively.

As in the two-class case, we can take a particular intersection of the likelihood ratio curve with a decision boundary, say  $LR_2 = 1$ , and find the corresponding value of  $x_1$  from Eq. 15; this value of  $x$  represents a threshold in the underlying data space, such that class  $\pi_2$  is decided for  $\mathbf{x} < x_1$ , and class  $\pi_3$  is decided for  $\mathbf{x} \geq x_1$ . Similarly,




Figure 5. The likelihood ratio curve when the underlying data are drawn from unit-variance normal distributions with parameters  $\mu_1 = 1$ ,  $\mu_2 = -1$ , and  $\mu_3 = 0$ , and where the decision boundary line segments are shown with dashed lines. An expression such as “ $\pi_1$ ” indicates a region where a particular decision is made.




Figure 6. The data space thresholds, shown as dashed lines, corresponding to the decision boundaries of Fig. 5, plotted against the underlying data pdfs ( $\mu_2 = -1$ ,  $\mu_3 = 0$ ,  $\mu_1 = 1$ , all variances equal to one). An expression such as “ $\pi_1$ ” indicates a region where a particular decision is made.

the decision boundary  $LR_1 = 1$  corresponds to a threshold  $x_2$  via Eq. 14 such that class  $\pi_3$  is decided for  $\mathbf{x} \leq x_2$  and class  $\pi_1$  is decided for  $\mathbf{x} > x_2$ . The resulting underlying data space representation of this decision strategy is exactly the same as that shown in Fig. 3.

Now consider a different choice of ideal observer decision criteria<sup>5,14</sup> such that the decision boundaries are  $LR_2 = LR_1 - 3/4$ ,  $LR_1 = 1$ , and  $LR_2 = -2LR_1 + 9/4$ , as shown in Fig. 5. Note that there are now three points at which the likelihood ratio curve and decision boundaries intersect. By finding the corresponding values of  $x$  for these intersection points, an underlying data space representation of this decision rule can again be found, as is shown in Fig. 6. Note that even in this simplest case (in terms of the underlying data PDFs, which we assumed to be equivariant), a possible behavior of the ideal observer is to partition the underlying data space into *four* regions with a set of *three* thresholds on  $x$ , an obvious increase in complexity when compared with the intuitively expected situation of Fig. 3.




Figure 7. The likelihood ratio curve when the underlying data are drawn from unit-variance normal distributions with parameters  $\mu_1 = 1/5$ ,  $\sigma_1^2 = 1/5$ ,  $\mu_2 = -3/4$ , and  $\sigma_2^2 = 1/2$ , and where the decision boundary line segments are shown with dashed lines. An expression such as “ $\pi_1$ ” indicates a region where a particular decision is made.




Figure 8. The data space thresholds, shown as dashed lines, corresponding to the decision boundaries of Fig. 7, plotted against the underlying data pdfs ( $\mu_1 = 1/5$ ,  $\sigma_1^2 = 1/5$ ,  $\mu_2 = -3/4$ , and  $\sigma_2^2 = 1/2$ ). An expression such as “ $\pi_1$ ” indicates a region where a particular decision is made.

The preceding examples were for the special case in which all three underlying data PDFs have the same variance. In general, of course, these variances may be different; in such a case, it can be shown that the likelihood ratio curve will form a closed loop in the likelihood ratio plane.<sup>11</sup> (Cases in which two of the three variances are equal, with the third differing, can also readily be treated explicitly,<sup>11</sup> but are beyond the scope of the present work.) One example of such a closed curve with underlying data parameters  $\mu_1 = 1/5$ ,  $\sigma_1^2 = 1/5$ ,  $\mu_2 = -3/4$ , and  $\sigma_2^2 = 1/2$ , and with choices of the five decision criteria such that the decision boundaries are  $LR_2 = LR_1 - 3/2$ ,  $LR_2 = -LR_1 + 5/2$ , and  $LR_2 = 1/2$ , is shown in Fig. 7. As in the example of Figs. 5 and 6, the decision boundaries intersect the likelihood ratio curve in three places, corresponding to three thresholds which partition the underlying data space into four regions, as shown in Fig. 8.

It is not difficult to show that the situation just described does not represent the upper bound on the complexity of the decision rule in its underlying data space representation. For the same set of parameters of




Figure 9. The likelihood ratio curve when the underlying data are drawn from unit-variance normal distributions with parameters  $\mu_1 = 1/5$ ,  $\sigma_1^2 = 1/5$ ,  $\mu_2 = -3/4$ , and  $\sigma_2^2 = 1/2$ , and where the decision boundary line segments are shown with dashed lines. An expression such as “ $\pi_1$ ” indicates a region where a particular decision is made.




Figure 10. The data space thresholds, shown as dashed lines, corresponding to the decision boundaries of Fig. 9, plotted against the underlying data pdfs ( $\mu_1 = 1/5$ ,  $\sigma_1^2 = 1/5$ ,  $\mu_2 = -3/4$ , and  $\sigma_2^2 = 1/2$ ). An expression such as “ $\pi_1$ ” indicates a region where a particular decision is made.

the PDFs of the underlying data ( $\mu_1 = 1/5$ ,  $\sigma_1^2 = 1/5$ ,  $\mu_2 = -3/4$ , and  $\sigma_2^2 = 1/2$ ), we can choose a different set of decision criteria such that the decision boundary line segments are now  $LR_2 = LR_1 - 3/2$ ,  $LR_1 = 2$ , and  $LR_2 = -7/8LR_1 + 9/4$ , as is shown in Fig. 9. There are now five intersection points of the likelihood ratio curve with the decision boundaries; each of these intersection points corresponds to a threshold in the underlying data space, which together partition the underlying data space into six regions, as shown in Fig. 10.

For our final example, we again consider the same set of underlying data parameters ( $\mu_1 = 1/5$ ,  $\sigma_1^2 = 1/5$ ,  $\mu_2 = -3/4$ , and  $\sigma_2^2 = 1/2$ ), and choose decision criteria such that the decision boundary line segments are  $LR_2 = 2LR_1 - 37/8$ ,  $LR_1 = 2$ , and  $LR_2 = -9/8LR_1 + 23/8$ ; this situation is illustrated in Fig. 11. There are now six intersection points of the likelihood ratio curve with the decision boundary line segments. Note that the intersection of the three decision boundaries lies outside the likelihood ratio loop; because of this, the slope,  $m_{13}$ , of the “1-vs.-3” decision boundary (here shown to be infinite) could actually take on any value such




Figure 11. The likelihood ratio curve when the underlying data are drawn from unit-variance normal distributions with parameters  $\mu_1 = 1/5$ ,  $\sigma_1^2 = 1/5$ ,  $\mu_2 = -3/4$ , and  $\sigma_2^2 = 1/2$ , and where the decision boundary line segments are shown with dashed lines. Because the intersection of the decision boundaries is outside the likelihood ratio curve, the “1-vs.-3” decision boundary could take on any slope in its allowed range (shown with dotted lines) without altering any of the intersection points of the likelihood ratio curve with the decision boundaries. An expression such as “ $\pi_1$ ” indicates a region where a particular decision is made.




Figure 12. The data space thresholds, shown as dashed lines, corresponding to the decision boundaries of Fig. 11, plotted against the underlying data pdfs ( $\mu_1 = 1/5$ ,  $\sigma_1^2 = 1/5$ ,  $\mu_2 = -3/4$ , and  $\sigma_2^2 = 1/2$ ). An expression such as “ $\pi_1$ ” indicates a region where a particular decision is made.

that  $-\infty < m_{13} \leq -9/8$  or  $2 \leq m_{13} < \infty$ ,<sup>14</sup> without altering any of the intersection points of the likelihood ratio curve with the decision boundaries. This is expressed by the dotted lines in Fig. 11. Since the intersection points would be unaffected by such a change in slope of the “1-vs.-3” line, the corresponding thresholds in the underlying data space, and hence the misclassification probabilities which form the coordinates of ROC space, are also independent of this slope. The decision strategy mapped into the underlying data space is shown in Fig. 12.

## 4. DISCUSSION

Even given the limited number of examples we have looked at here, it is already clear that the behavior of the three-class ideal observer for univariate underlying data can be considerably more complex than would be expected from the model shown in Fig. 3, a straightforward generalization of the two-class model of Fig. 1. For the univariate trinormal model, a more complete “catalog” of the possible range of behaviors (the shapes of the likelihood ratio curve, and the resulting number of intersections of that curve with the decision boundaries) can be made.<sup>11</sup> The overall conclusion is found to be essentially the same: the ideal observer’s decision strategy, when represented in the underlying data space, will partition that space with as many as six decision thresholds.

Based on these results, we have also been able to generalize the expressions of Eqs. 7 and 8 to the univariate trinormal case,<sup>12</sup> deriving expressions for the misclassification probabilities in terms of the decision thresholds in the underlying data space (which in turn are derived from the intersections of the likelihood ratio curve with the decision boundary line segments). When the resulting expressions for the misclassification probabilities—*i.e.*, the observer’s ROC operating point coordinates—are compared with the observations already obtained<sup>11</sup> (the most salient of which are presented in Sec. 3 above) regarding the intersections of the likelihood ratio curve and the decision boundaries, it is even possible to draw conclusions about the dimensionality of the resulting ROC hypersurface.<sup>12</sup>

For example, in the particular situation illustrated in Fig. 11, we pointed out that the intersection points of the likelihood ratio curve with the decision boundaries cannot depend on the value of the slope of the “1-*vs.*-3” decision boundary (which lies entirely outside the likelihood ratio curve). This in turn means that the observer’s ROC operating point coordinates (the misclassification probabilities) cannot depend on this slope at that particular operating point. This slope is one of five quantities (the two coordinates of the intersection of the decision boundaries, and the three slopes of those boundaries) which can be mapped in a one-to-one fashion to the set of five decision criteria used by the three-class ideal observer to make decisions.<sup>14</sup> It follows that the observer’s ROC hypersurface at this operating point can have at most four degrees of freedom: each of the six misclassification probabilities is, in general, a function of the five decision criteria; but if, at a given point, they only depend on four (or fewer) of those decision criteria, then that is the upper bound on the dimensionality of the hypersurface at that operating point.

This result, particularly if generalizable, would have negative implications for our proposed three-class performance metric,<sup>2</sup> as mentioned in the Introduction. We consider it to be of great importance, however, for two reasons. First, it is crucial to develop an accurate understanding of the actual behavior of the ROC hypersurface of a given observer; a successful performance metric must be made to suit this behavior. Second, although the present results were obtained only for univariate trinormal data, we feel the insights we have gained through explicit analysis of this case will prove invaluable in attacking more general situations, which in turn should allow us to either refine our proposed performance metric or develop one more suitable to the actual properties of the ROC hypersurface of the three-class ideal observer.

## 5. CONCLUSIONS

We have explicitly analyzed the decision strategy of the three-class ideal observer for univariate trinormal underlying data. We found that the representation of this decision strategy in the underlying data space is more complicated than might intuitively be expected, involving as many as six decision thresholds rather than just two. This underlying data space representation, however, can facilitate the explicit calculation of operating points on the observer’s ROC hypersurface. We expect that the insights gained from this work will allow us to pursue more sophisticated data models, and thus develop a suitable performance metric for evaluating observer performance in a three-class classification task.

## ACKNOWLEDGMENTS

The authors would like to thank Craig Abbey, whose questions and comments on this subject inspired us to investigate it more thoroughly than we had previously.

This work was supported by grant R01 EB000863-05 from the National Institutes of Health (Kevin Berbaum, principal investigator) through a subcontract from the University of Iowa to the University of Chicago (Charles Metz, subcontract principal investigator).

## REFERENCES

- [1] Edwards, D. C., Metz, C. E., and Nishikawa, R. M., “The hypervolume under the ROC hypersurface of ‘near-guessing’ and ‘near-perfect’ observers in  $N$ -class classification tasks,” *IEEE Trans. Med. Imag.* **24**, 293–299 (2005).
- [2] Edwards, D. C. and Metz, C. E., “A utility-based performance metric for ROC analysis of  $N$ -class classification tasks,” in [Proc. SPIE Vol. 6515 *Medical Imaging 2007: Image Perception, Observer Performance, and Technology Assessment*], Yulei Jiang and Berkman Sahiner, eds., 6515031–65150310 (2007).
- [3] Edwards, D. C. and Metz, C. E., “Optimality of a utility-based performance metric for ROC analysis,” in [Proc. SPIE Vol. 6917 *Medical Imaging 2008: Image Perception, Observer Performance, and Technology Assessment*], Berkman Sahiner and David J. Manning, eds., 69170F1–68170F6 (2008).
- [4] Edwards, D. C. and Metz, C. E., “Comparing the performance of two observers using a novel utility-based performance metric for ROC analysis,” in [Proc. SPIE Vol. 7263 *Medical Imaging 2009: Image Perception, Observer Performance, and Technology Assessment*], Berkman Sahiner and David J. Manning, eds., 72630W1–72630W7 (2009).
- [5] Edwards, D. C., Metz, C. E., and Kupinski, M. A., “Ideal observers and optimal ROC hypersurfaces in  $N$ -class classification,” *IEEE Trans. Med. Imag.* **23**, 891–895 (2004).
- [6] He, X. and Frey, E., “Revisiting the general three-class ROC theory: A decision theoretic approach.” Medical Image Perception Conference XIII, Santa Barbara, CA (2009).
- [7] Edwards, D. C. and Metz, C. E., “Optimization of restricted ROC surfaces in three-class classification tasks,” *IEEE Trans. Med. Imag.* **26**, 1345–1356 (2007).
- [8] Scurfield, B. K., “Multiple-event forced-choice tasks in the theory of signal detectability,” *J. Math. Psychol.* **40**, 253–269 (1996).
- [9] Metz, C. E., Herman, B. A., and Shen, J.-H., “Maximum-likelihood estimation of receiver operating characteristic (ROC) curves from continuously-distributed data,” *Statist. Med.* **17**, 1033–1053 (1998).
- [10] Metz, C. E. and Pan, X., “‘Proper’ binormal ROC curves: Theory and maximum-likelihood estimation,” *J. Math. Psychol.* **43**, 1–33 (1999).
- [11] Edwards, D. C. and Metz, C. E., “The three-class ideal observer for univariate normal data: Decision variable and threshold properties,” *IEEE Trans. Med. Imag.* (2009). In review.
- [12] Edwards, D. C. and Metz, C. E., “The three-class ideal observer for univariate normal data: ROC surface properties,” *IEEE Trans. Med. Imag.* (2010). In review.
- [13] Papoulis, A., [*Probability, Random Variables, and Stochastic Processes*], McGraw-Hill, Inc., New York (1991).
- [14] Edwards, D. C. and Metz, C. E., “Restrictions on the three-class ideal observer’s decision boundary lines,” *IEEE Trans. Med. Imag.* **24**, 1566–1573 (2005).