Evaluating Classification Performance:

Receiver Operating Characteristic and Expected Utility

Supplemental Materials

Contents

Operating Criterion for Maximizing Expected Utility

For a binary classification, expected utility can be estimated,

***Equation 1*. Expected Utility of Classification**

$$EU = p_1[P(1|1)u(1|1) + P(0|1)u(0|1)] + p_0[P(1|0)u(1|0) + P(0|0)u(0|0)]$$

How should a classifier maximize its expected utility? For any given prior probabilities and utilities, Equation 2 specifies the condition under which a classifier will maximize expected utility (Green & Swets, 1966; Johnson & Wichern, 2007).

***Equation 2*. Operating Criterion for Maximizing Expected Utility**

$$L(\gamma) = \frac{f_1(\gamma)}{f_0(\gamma)} = \left(\frac{p_0}{p_1}\right)\left(\frac{u(0|0) - u(1|0)}{u(1|1) - u(0|1)}\right) = \frac{r}{o}$$

$$\underset{\textit{likelihood ratio}}{\Uparrow} \qquad = \qquad \underset{\textit{utility-prior ratio}}{\Uparrow}$$

The above condition can be derived from the expected utility equation (Equation 1).

$$EU = p_1[P(1|1)u(1|1) + P(0|1)u(0|1)] + p_0[P(1|0)u(1|0) + P(0|0)u(0|0)]$$

$$= p_1P(1|1)u(1|1) + p_1[1 - P(1|1)]u(0|1) + p_0P(1|0)u(1|0)$$

$$+ p_0[1 - P(1|0)]u(0|0)$$

$$= p_1u(0|1) + p_0u(0|0) + p_1[u(1|1) - u(0|1)]P(1|1)$$

$$+ p_0[u(1|0) - u(0|0)]P(1|0)$$

$$= p_1u(0|1) + p_0u(0|0) + p_1[u(1|1) - u(0|1)][1 - F_1(\gamma)]$$

$$+ p_0[u(1|0) - u(0|0)][1 - F_0(\gamma)]$$

Assuming the prior probabilities and utilities are constants, we want to find out the response criterion that can optimize expected utility,

$$\frac{\partial EU}{\partial \gamma} = 0$$

$$\frac{\partial EU}{\partial \gamma} = -p_1[u(1|1) - u(0|1)]\frac{\partial [F_1(\gamma)]}{\partial \gamma} - p_0[u(1|0) - u(0|0)]\frac{\partial [F_0(\gamma)]}{\partial \gamma}$$

$$= -p_1[u(1|1) - u(0|1)]f_1(\gamma) + p_0[u(0|0) - u(1|0)]f_0(\gamma) = 0$$

$$\Longrightarrow L(\gamma) = \frac{f_1(\gamma)}{f_0(\gamma)} = \left(\frac{p_0}{p_1}\right)\left(\frac{u(0|0) - u(1|0)}{u(1|1) - u(0|1)}\right) = \frac{r}{o}$$

Expected Utility Line in an ROC space

What is the relation between ROC analysis and expected utility analysis? One can visualize a classifier's expected utility in its ROC space for any given prior probabilities and utilities (Provost & Fawcett, 2001). The expected utility equation (Equation 1) is equivalent to a line in the ROC space, as shown in Equation 3 below.

***Equation 3*. Expected Utility Line in an ROC space**

$$P(1|1) = \frac{EU - w}{v} + \frac{r}{o}P(1|0) = a + bP(1|0)$$

$$v = p_1[u(1|1) - u(0|1)]$$

$$w = p_0 u(0|0) + p_1 u(0|1)$$

$$Intercept \equiv a = \frac{EU - w}{v}$$

$$Slope \equiv b = \frac{r}{o}$$

Equation 3 is simply a transformation of Equation 1,

$$EU = p_1[P(1|1)u(1|1) + P(0|1)u(0|1)] + p_0[P(1|0)u(1|0) + P(0|0)u(0|0)]$$

$$= p_1 P(1|1)u(1|1) + p_1[1 - P(1|1)]u(0|1) + p_0 P(1|0)u(1|0)$$

$$+ p_0[1 - P(1|0)]u(0|0)$$

$$= p_1 u(0|1) + p_0 u(0|0) + p_1[u(1|1) - u(0|1)]P(1|1)$$

$$+ p_0[u(1|0) - u(0|0)]P(1|0)$$

We can move $P(1|1)$ to the left side of the equation and everything else to the right side,

$$p_1[u(1|1) - u(0|1)]P(1|1) = EU - [p_1 u(0|1) + p_0 u(0|0)] - p_0[u(1|0) - u(0|0)]P(1|0)$$

$$\Longrightarrow P(1|1) = \frac{EU - [p_1 u(0|1) + p_0 u(0|0)] - p_0[u(1|0) - u(0|0)]P(1|0)}{p_1[u(1|1) - u(0|1)]}$$

$$= \frac{EU - [p_1 u(0|1) + p_0 u(0|0)]}{p_1[u(1|1) - u(0|1)]} - \frac{p_0[u(1|0) - u(0|0)]}{p_1[u(1|1) - u(0|1)]} P(1|0)$$

$$= \frac{EU - w}{v} - \frac{r}{o} P(1|0)$$

$$v = p_1[u(1|1) - u(0|1)]$$

$$w = p_1 u(0|1) + p_0 u(0|0)$$

$$\frac{r}{o} = \frac{p_0[u(1|0) - u(0|0)]}{p_1[u(1|1) - u(0|1)]}$$

Utilities of Investigator Classification Outcomes

Investigator classification outcomes can lead to serious legal consequences. For example, if investigators falsely classify a guilty suspect as innocent, the guilty suspect may evade punishment and continue to harm society. If investigators falsely classify an innocent suspect as guilty, not only will the guilty suspect evade punishment, but the innocent suspect will suffer as a result. Although it may be difficult to estimate the absolute values of these consequences, it is less difficult to estimate their ratio (Yang et al., 2019).

To estimate the utilities of investigator classification outcomes, let us first define a reference point. Here I choose the goal of the police investigation as the reference point (other choices of the reference point does not change the utility ratio; see Yang et al., 2019). Thus, the utility of a classification outcome is determined by the *discrepancy* between this reference point and the consequences associated with the classification outcome.

According to the Model Code of Criminal Procedure (O'Connor & Rausch, 2008), one of the primary goals of a criminal investigation is to "establish, with regard to a specific criminal offense, if a suspect can be identified; and […] determine whether sufficient reasons exist for the prosecution of a suspect of a criminal offense" (p. 154). Thus, I consider it reasonable to assume that a successful police investigation for a criminal offense is one that identifies the culprit and documents evidence establishing the guilt of that person. If a classification outcome successfully achieves the goal of identifying and incriminating the culprit, I assume the utility associated with this outcome to be zero. However, if a classification outcome does not achieve this goal, I assume the utility is negative (or equivalently, a cost occurs). The greater the discrepancy between the outcome and the investigation goal, the smaller the utility (i.e., the higher the cost).

With the reference point, let us examine each of the four classification outcomes.

Outcome $G|G$. When a guilty suspect is correctly classified as guilty, police achieve their goal of identifying and incriminating the culprit. Therefore, the utility is zero, $u(G|G) = 0$.

Outcome $I|G$. When a guilty suspect is falsely classified as innocent, police have not achieved their goal of identifying and incriminating the culprit. Accordingly, the utility is negative, $u(I|G) = u(failing\ to\ incriminate\ a\ guilty)$.

Outcome $G|I$. When an innocent suspect is mistakenly classified as guilty, police have not only failed to achieve their goal of identifying the culprit, but they have also incriminated the wrong person. Accordingly, the utility associated with the false classification of an innocent suspect includes both the utility of failing to incriminate the guilty suspect and the utility of wrongfully incriminating an innocent suspect, $u(G|I) = u(failing\ to\ incriminate\ a\ guilty) + u(incriminating\ an\ innocent)$.

Outcome $I|I$. The final outcome pertains to the scenario in which an innocent suspect is not identified. Some researchers have assumed no cost for this outcome because it is a correct classification (e.g., Clark, 2012; Malpass, 2006). However, considering the reference point, a discrepancy still exists between this outcome and the police goal: police have not identified and incriminated the culprit even though the innocent suspect has not been classified as guilty. Therefore, the utility of this outcome is negative, $u(I|I) = u(failing\ to\ incrimimate\ a\ guilty)$

Thus, the utility ratio of investigator classification outcomes can be reduced to a ratio between two components.

$$r = \frac{u(I|I) - u(G|I)}{u(G|G) - u(I|G)} = \frac{u(incriminating\ an\ innocent)}{u(failing\ to\ incriminate\ a\ guilty)}$$

The most commonly referenced version of this ratio was introduced by Blackstone (1769): "For the law holds, that it is better that ten guilty persons escape, than that one innocent suffer" (as cited in Volokh, 1997). Therefore, the example assumes a utility ratio of 10.

Alternative Selection Methods

Alternative to EU lines, there exist methods that select the optimal ROC point without explicitly specifying prior probabilities and utilities. These methods include the Yonden index (Youden, 1950), the closest-to-(0, 1) criterion (Perkins & Schisterman, 2006), and the concordance probability (Liu, 2012). Yet, rather than avoiding the trouble to specify prior probabilities and utilities, these methods implicitly assume values for these quantities.

Youden Index (J)

For example, the Youden index (*J*) selects the optimal ROC point that can maximize the difference between true positive and false positive rates, $J = \max\{true\ positve - false\ positve\}$ (Perkins & Schisterman, 2006). Accordingly, one can derive the condition to optimize *J*,

$$\frac{\partial J}{\partial \gamma} = \frac{\partial[P(1|1) - P(1|0)]}{\partial \gamma} = \frac{\partial[\left(1 - F_1(\gamma)\right) - (1 - F_0(\gamma))]}{\partial \gamma} = f_0(\gamma) - f_1(\gamma) = 0$$

$$\Longrightarrow L(\gamma) = \frac{f_1(\gamma)}{f_0(\gamma)} = 1$$

Geometrically, this is equivalent to find an ROC point of which the tangent line has a slope of 1. When the utility-prior ratio is equal to the likelihood ratio of 1, the selected ROC point maximizes expected utility. In other words, maximizing Youden index is equivalent to locating the MEU line with a slope of 1.

Closest-to-(0, 1) Criterion

The closest-to-(0, 1) criterion selects the optimal ROC point that is the closest to the perfection point—the (0, 1) point in an ROC space (Perkins & Schisterman, 2006). Accordingly, one can derive the condition to optimize the distance between an ROC curve and the (0, 1) point,

$$\frac{\partial[\left(0-P(0|1)\right)^2+\left(1-P(1|1)\right)^2]}{\partial\gamma}=\frac{\partial[(F_0(\gamma)-1)^2+F_1^2(\gamma)]}{\partial\gamma}$$

$$=2(F_0(\gamma)-1)f_0(\gamma)+2F_1(\gamma)f_1(\gamma)=0$$

$$\Longrightarrow L(\gamma)=\frac{f_1(\gamma)}{f_0(\gamma)}=\frac{1-F_0(\gamma)}{F_1(\gamma)}=\frac{P(1|0)}{1-P(1|1)}$$

Geometrically, this is equivalent to find an ROC point of which the tangent line is perpendicular to its connection line with the (0, 1) point. When the utility-prior ratio is equal to the likelihood ratio of $\frac{false\ positive}{1-true\ positive}$, the selected ROC point maximizes expected utility. In other words, the closest-to-(0, 1) criterion assumes a utility-prior ratio of $\frac{false\ positive}{1-true\ positive}$ for the selected ROC point to maximize expected utility.

Concordance Probability

The concordance probability criterion selects the optimal ROC point that can maximize the concordance probability of a dichotomized measure, or geometrically, the area of a rectangle associated with the ROC point (Liu, 2012). Accordingly, one can derive the condition to optimize the concordance probability,

$$\frac{\partial[P(1|1)(1-P(1|0)]}{\partial\gamma}=\frac{\partial[(1-F_1(\gamma))F_0(\gamma)]}{\partial\gamma}=\left(1-F_1(\gamma)\right)f_0(\gamma)-F_0(\gamma)f_1(\gamma)=0$$

$$\Longrightarrow L(\gamma)=\frac{f_1(\gamma)}{f_0(\gamma)}=\frac{1-F_1(\gamma)}{F_0(\gamma)}=\frac{P(1|1)}{1-P(1|0)}$$

Geometrically, this is equivalent to find an ROC point of which the tangent line is parallel to the diagonal of the rectangle associated with the ROC point. When the utility-prior ratio is equal to the likelihood ratio of $\frac{true\ positive}{1-false\ positive}$, the selected ROC point maximizes expected utility. In other words, the concordance probability criterion assumes a utility-prior ratio of $\frac{true\ positive}{1-false\ positive}$ for the selected ROC point to maximize expected utility.

References

Clark, S. E. (2012). Costs and benefits of eyewitness identification reform: Psychological science and public policy. *Perspectives on Psychological Science, 7*, 238-259. doi:10.1177/1745691612439584

Malpass, R. S. (2006). A policy evaluation of simultaneous and sequential lineups. *Psychology, Public Policy, and Law*, *12*, 394-418. doi:10.1037/1076-897.12.4.394

O'Connor, V., & Rausch, C. (Eds.). (2008). *Model codes for post-conflict criminal justice* (Vol. 2). United States Institute of Peace Press.

Other references are included in the reference list of the manuscript.