STATISTICS IN MEDICINE
Statist. Med. 2004; **23**:2257–2266 (DOI: 10.1002/sim.1835)

Advantages to transforming the receiver operating characteristic (ROC) curve into likelihood ratio co-ordinates

Nils P. Johnson*,†

College of Physicians and Surgeons, Columbia University, 630 W. 168th St., New York, NY 10032, U.S.A.

SUMMARY

Traditionally, the receiver operating characteristic (ROC) curve for a diagnostic test plots true positives (sensitivity) against false positives (one minus specificity). However, this representation brings with it several drawbacks. A transformation to positive and negative likelihood ratio co-ordinates, scaled by base-ten logarithms, offers several advantages. First we motivate the use of positive and negative likelihood ratios, emphasizing their relationship to modification of the odds ratio. Then we highlight properties of likelihood ratios using the traditional ROC axes. Finally, we demonstrate ROC curves and their properties after conversion to likelihood ratio co-ordinates. These graphs do not waste space for tests lacking diagnostic power, and offer a simple visual assessment of a test's impact on the odds ratio.

KEY WORDS: receiver operating characteristic (ROC) curve; likelihood ratio; sensitivity; specificity

INTRODUCTION

A diagnostic test revises the pre-test probability of disease via application of Bayes' theorem [1]. For a binary diagnostic test, a positive result revises this probability upwards while a negative result revises it downwards. To convert a continuous diagnostic test to a binary one, ranges of results in the continuous variable must be interpreted as either positive or negative. The choice of positive and negative cut-offs determines the rates of true and false positives, when compared to the gold standard for diagnosing the disease. The true and false positive rates, in turn, determine the magnitude of revision to the disease's probability.

A receiver operating characteristic (ROC) curve plots the true positive rate against the false positive rate for different choices of cut-off ranges [2]. However, these axes obscure more than they enlighten. First, half of the area shown on an ROC plot cannot contain any meaningful curves. This deceives the eye into minimizing area differences. Second, it requires substantial familiarity with the subject to understand how the true positive or false positive rate alters the probability of disease. Positive and negative likelihood ratios provide clearer descriptors

*Correspondence to: Nils P. Johnson, 100 Haven Avenue, #10A, New York, NY 10032, U.S.A.
†E-mail: nj90@columbia.edu

Received June 2003
Accepted February 2004

for altering disease probability. Third, several constraints and properties of ROC curves can be more easily represented when viewed using positive and negative likelihood ratio axes.

We propose that likelihood ratios provide a better set of axes for the ROC curve than the traditional true and false positive rates. First, we motivate the definition of positive and negative likelihood ratios. Their modification of the odds ratio takes on particular importance. Second, we note some important properties of ROC curves. Third, we graphically demonstrate several connections between likelihood ratios and the traditional ROC axes. Finally, we demonstrate the appearance and properties of the ROC curve with positive and negative likelihood ratio axes, scaled by base-ten logarithms. These ease visual interpretation for selecting an operating point for a binary diagnostic test.

LIKELIHOOD RATIOS

To motivate the new axes for ROC curves, let us derive some well-known theory concerning diagnostic tests. Four outcomes can result when comparing a binary diagnostic test T to the gold-standard definition of disease D. Either a positive, T_+, or negative, T_-, test result can exist in the presence, D_+, or absence, D_-, of disease. The sensitivity, $\mathsf{Se} = P(T_+|D_+)$ (probability of a positive test result in the presence of disease), and specificity, $\mathsf{Sp} = P(T_-|D_-)$ (probability of a negative test result in the absence of disease), characterize a binary diagnostic test. A pre-test probability $p = P(D_+)$, for example the prevalence of disease in the general population [1], gives the following two-by-two table of outcome probabilities.

	D_+	D_-
T_+	$p\,\mathsf{Se}$	$(1-p)(1-\mathsf{Sp})$
T_-	$p(1-\mathsf{Se})$	$(1-p)\mathsf{Sp}$
Total	p	$1-p$

To compute $p_+ = P(D_+|T_+)$, the probability of disease given a positive test result, apply Bayes' theorem [3] as guided by this outcome table to find

$$p_+ = \frac{P(D_+)P(T_+|D_+)}{P(D_+)P(T_+|D_+) + P(D_-)P(T_+|D_-)} = \frac{p\,\mathsf{Se}}{p\,\mathsf{Se} + (1-p)(1-\mathsf{Sp})}$$

Manipulation gives the more interpretable form

$$\frac{p_+}{1-p_+} = \frac{\mathsf{Se}}{1-\mathsf{Sp}}\,\frac{p}{1-p} \tag{1}$$

The odds ratio $\mathsf{OR} = p/(1-p)$ describes the ratio of disease probability to the probability of no disease. An odds ratio takes on values between zero (certainly no disease present) and infinity (certainly disease present) and can be converted back to a disease probability by $p = \mathsf{OR}/(1+\mathsf{OR})$. The expression

$$\mathsf{LR}_+ = \frac{\mathsf{Se}}{1-\mathsf{Sp}} \tag{2}$$

is called the positive likelihood ratio, or the likelihood ratio of a positive test. The definitions of the odds ratio and positive likelihood ratio allow equation (1) to be written as the product

$$\mathsf{OR}_+ = \mathsf{LR}_+ \cdot \mathsf{OR} \tag{3}$$

Similarly, to compute $p_- = P(D_+|T_-)$, the probability of disease despite a negative test result, apply Bayes' theorem as above to find

$$p_- = \frac{P(D_+)P(T_-|D_+)}{P(D_+)P(T_-|D_+) + P(D_-)P(T_-|D_-)} = \frac{p(1-\mathsf{Se})}{p(1-\mathsf{Se}) + (1-p)\mathsf{Sp}}$$

Manipulation as before yields

$$\frac{p_-}{1-p_-} = \frac{1-\mathsf{Se}}{\mathsf{Sp}} \frac{p}{1-p} \tag{4}$$

where the expression

$$\mathsf{LR}_- = \frac{1-\mathsf{Se}}{\mathsf{Sp}} \tag{5}$$

is called the negative likelihood ratio, or the likelihood ratio of a negative test. The definitions of the odds ratio and negative likelihood ratio allow equation (4) to be written as the product

$$\mathsf{OR}_- = \mathsf{LR}_- \mathsf{OR} \tag{6}$$

A useful diagnostic test must revise the pre-test probability upwards after a positive test result, or revise the pre-test probability downwards after a negative test result. Equations (3) and (6) state that, for this to hold, $\mathsf{LR}_+ > 1$ and $\mathsf{LR}_- < 1$, as the odds ratio changes in the same direction as the probability. Both of these conditions yield the same constraint: $\mathsf{Se} + \mathsf{Sp} > 1$. Fortunately this constraint can always hold, for if $\mathsf{Se}+\mathsf{Sp} < 1$ then reversing T_+ and T_- creates a new diagnostic test satisfying the constraint, as can be seen by examination of the two-by-two table of outcome probabilities. The case $\mathsf{Se} + \mathsf{Sp} = 1$ describes a diagnostic test equivalent to a coin toss with probability Se of T_+, and offers no diagnostic power as $\mathsf{LR}_+ = \mathsf{LR}_- = 1$.

Likelihood ratios provide an alternate description to sensitivity and specificity for a binary diagnostic test. Their usefulness can be seen best in equations (3) and (6), where they revise the pre-test odds ratio to give the post-test odds ratio in a multiplicative fashion. This post-test odds ratio serves, in turn, as the pre-test odds ratio for the next independent diagnostic test, thereby concatenating equations of form (3) and (6). After performing a series of independent diagnostic tests on a patient with an initial odds ratio OR_0, the final odds ratio OR_{a+b} can be written as a product

$$\mathsf{OR}_{a+b} = \mathsf{LR}_+^{(1)} \cdots \mathsf{LR}_+^{(a)} \cdot \mathsf{LR}_-^{(1)} \cdots \mathsf{LR}_-^{(b)} \cdot \mathsf{OR}_0$$

where $a + b$ diagnostic tests have yielded a positive results and b negative results, with likelihood ratios $\mathsf{LR}_{+/-}^{(i)}$ for each test. Dividing by the original odds ratio and taking the logarithm of both sides yields

$$\log \frac{\mathsf{OR}_{a+b}}{\mathsf{OR}_0} = \sum_{i=1}^{a} \log \mathsf{LR}_+^{(i)} + \sum_{j=1}^{b} \log \mathsf{LR}_-^{(j)} \tag{7}$$

where the first summation describes the net increase to the odds ratio provided by positive test results and the second summation describes the net decrease to the odds ratio provided by negative test results. Since $\mathsf{LR}_+ > 1$, then $\log \mathsf{LR}_+ > 0$ and the first summation is strictly positive. Since $\mathsf{LR}_- < 1$, then $\log \mathsf{LR}_- < 0$ and the second summation is strictly negative.

Equation (7) provides a compact form for summarizing the effect of multiple independent diagnostic test results on the probability of disease. For example, to increase the odds ratio by at least an order of magnitude from its initial value, $\mathsf{OR}_{a+b}/\mathsf{OR}_0 \geqslant 10$, a diagnostic test must have $\log_{10} \mathsf{LR}_+ \geqslant 1$. Similarly, to decrease the odds ratio by at least an order of magnitude from its initial value, a diagnostic test must have $\log_{10} \mathsf{LR}_- \leqslant -1$.

THE ROC CURVE

While a binary diagnostic test produces only two outcomes, a continuous diagnostic test returns a result from a continuum. A continuous test can be converted to a binary test by choosing ranges of the test outcome and interpreting results within those ranges as negative and results outside those ranges as positive. Ranges based on the likelihood-ratio criterion or its equivalent yield optimal results for a number of different decision goals [3].

Denote the probability density functions of a continuous diagnostic test in variable x as $f(x|D_+)$ in the presence of disease and $f(x|D_-)$ in the absence of disease. Compute the likelihood ratio

$$L(x) = \frac{f(x|D_+)}{f(x|D_-)} \tag{8}$$

which describes the ratio of the probability of observing x in the diseased population compared to the disease-free population. (Although equation (8) is called the likelihood ratio, it is a different concept than the likelihood ratios in equations (2) and (5). The term 'likelihood' is popular and appears in these distinct, but related, concepts.) Label regions with $L(x) < x_0$ as T_- and those with $L(x) \geqslant x_0$ as T_+, where x_0 can vary between zero and infinity. Each x_0 converts the continuous diagnostic test into a binary diagnostic test, with

$$\mathsf{Se} = P(T_+|D_+) = \int_{x_0}^{\infty} f(x|D_+)\,\mathrm{d}x$$

true positives and

$$1 - \mathsf{Sp} = P(T_+|D_-) = \int_{x_0}^{\infty} f(x|D_-)\,\mathrm{d}x$$

false positives, assuming that $L(x)$ increases monotonically with x. (If this is not the case, Se and $1 - \mathsf{Sp}$ can be expressed as a summation of integrals over appropriate regions as determined by $L(x)$. The following results hold in this case too [4].)

An ROC curve (also termed a relative operating characteristic curve by some authors [5]) plots Se (true positives) against $1-\mathsf{Sp}$ (false positives) for all possible likelihood ratio cut-offs. The slope of the ROC curve can be found by application of the chain rule, the fundamental

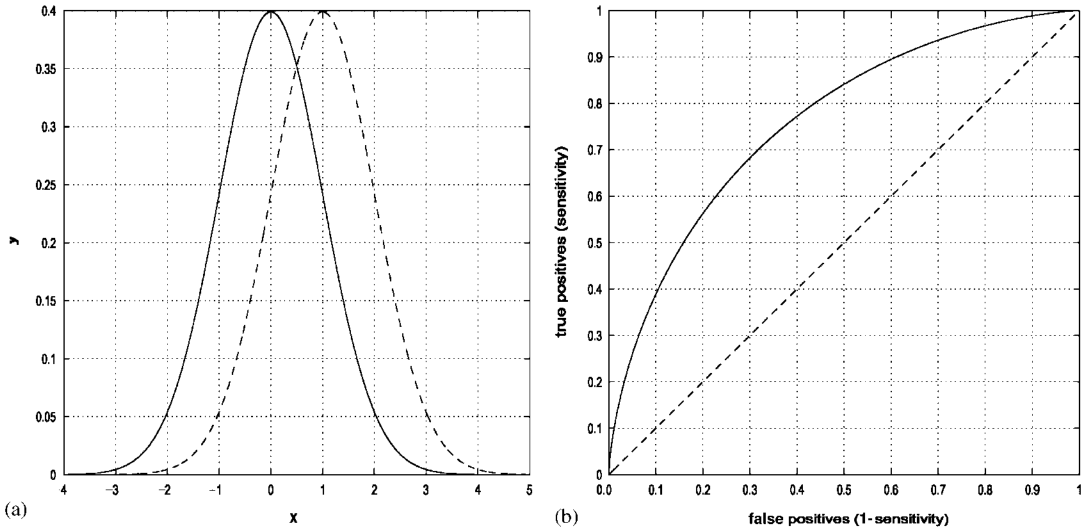

Figure 1. (a) The probability density function of a hypothetical continuous diagnostic test for non-diseased (solid) and diseased (dashed) populations. Each normal curve has a standard deviation of one and a mean of zero (solid) or one (dashed) in arbitrary units. (b) The ROC curve (solid) for the hypothetical continuous diagnostic test, as well as the line $\mathsf{Se} + \mathsf{Sp} = 1$ (dashed).

theorem of calculus, and by noting that $f(\infty|D_+) = f(\infty|D_-) = 0$,

$$\mathrm{d}\left(\frac{P(T_+|D_+)}{P(T_+|D_-)}\right)_{x_0} = \frac{\mathrm{d}P(T_+|D_+)/\mathrm{d}x_0}{\mathrm{d}P(T_+|D_-)/\mathrm{d}x_0} = \frac{f(x_0|D_+)}{f(x_0|D_-)} = L(x_0)$$

The slope of the ROC curve equals the likelihood ratio at that point as given by equation (8).

Figure 1 demonstrates a continuous diagnostic test with normal probability density functions separated by one standard deviation for the diseased and non-diseased populations, and the corresponding ROC curve. For a low cut-off value, almost all patients receive positive test results. This produces a high sensitivity, but a low specificity. For a high cut-off value, almost all patients receive negative test results. This produces a low sensitivity, but a high specificity. As the cut-off increases in Figure 1(a) from minus infinity to positive infinity, values along the ROC curve move from the upper-right to the lower-left corners of Figure 1(b).

An ROC curve produced using the likelihood-ratio criterion has several important properties [3, 4]. First, as demonstrated above, the slope equals the likelihood ratio at that point. Second, the slope decreases monotonically, which implies that the ROC curve cannot be concave upward. Third, the ROC curve is 'proper' in the sense that it defines the upper boundary of all possible ROC curves based on the underlying probability density functions, and lies entirely outside of the area $\mathsf{Se} + \mathsf{Sp} < 1$. Fourth, the area under the ROC curve (AUC) equals the percentage correct in a two-alternative forced-choice task [6]. However, the AUC has been criticized for being an inconsistent metric when comparing two curves [7].

LIKELIHOOD RATIOS ON THE ROC CURVE

Equation (2) can be manipulated into

$$\mathsf{Se} = \mathsf{LR}_{+}(1 - \mathsf{Sp}) \tag{9}$$

and equation (5) can be rewritten as

$$\mathsf{Se} = \mathsf{LR}_{-}(1 - \mathsf{Sp}) + (1 - \mathsf{LR}_{-}) \tag{10}$$

Lines through the lower-left corner ($\mathsf{Se}=0$ and $1-\mathsf{Sp}=0$) of an ROC curve have slope LR_{+}, as shown in equation (9). Lines through the upper-right corner ($\mathsf{Se}=1$ and $1-\mathsf{Sp}=1$) of an ROC curve have slope LR_{-}, as shown in equation (10).

As can be seen from equation (7), regions of the ROC curve where $|\log \mathsf{LR}_{+}| > |\log \mathsf{LR}_{-}|$ will proportionally increase the odds ratio more after a positive test result than its decrease after a negative test result. The equivalent statement $\mathsf{LR}_{+} > 1/\mathsf{LR}_{-}$ can locate these regions in the potential ROC space, using equations (2) and (5). Set $\mathsf{LR}_{+} = 1/\mathsf{LR}_{-}$ and cross multiply to find $\mathsf{Se}(1-\mathsf{Se}) = \mathsf{Sp}(1-\mathsf{Sp})$, which can be solved as a quadratic equation $\mathsf{Se}^2 - \mathsf{Se} + (\mathsf{Sp} - \mathsf{Sp}^2) = 0$ with Sp as the independent variable. The solution $\mathsf{Se} = \frac{1}{2} \mp \frac{1}{2} \pm \mathsf{Sp}$ yields the straight lines $\mathsf{Se} = \mathsf{Sp}$ and $\mathsf{Se} = 1 - \mathsf{Sp}$. A similar analysis finds that regions between these lines have $|\log \mathsf{LR}_{+}| > |\log \mathsf{LR}_{-}|$, while regions outside have this inequality reversed.

Figure 2 demonstrates several properties of positive and negative likelihood ratios in the potential area for an ROC curve. Three regions can be defined using the lines $\mathsf{Se} = \mathsf{Sp}$ and $\mathsf{Se} = 1 - \mathsf{Sp}$. The region below $\mathsf{Se} = 1 - \mathsf{Sp}$ has $\mathsf{LR}_{+} < 1$ and $\mathsf{LR}_{-} > 1$ and therefore no proper ROC curve will exist there. The region above $\mathsf{Se} = 1 - \mathsf{Sp}$ but below $\mathsf{Se} = \mathsf{Sp}$ has $|\log \mathsf{LR}_{+}| > |\log \mathsf{LR}_{-}|$, and operating points in this region produce proportionally larger increases after positive test results than negative test results. The region above $\mathsf{Se} = 1 - \mathsf{Sp}$ and $\mathsf{Se} = \mathsf{Sp}$ has $|\log \mathsf{LR}_{-}| > |\log \mathsf{LR}_{+}|$, and operating points in this region produce proportionally larger increases after negative test results than positive test results. Only along $\mathsf{Se} = \mathsf{Sp}$ and $\mathsf{Se} = 1 - \mathsf{Sp}$ does $|\log \mathsf{LR}_{+}| = |\log \mathsf{LR}_{-}|$. Straight lines through the lower-left corner represent isocontours of LR_{+} as given by equation (9). Straight lines through the upper-right corner represent isocontours of LR_{-} as given by equation (10).

THE ROC CURVE IN LIKELIHOOD RATIO CO-ORDINATES

Half of the potential region for ROC curves has no meaning. Proper ROC curves lie above $\mathsf{Se} = 1 - \mathsf{Sp}$, except at the lower-left and upper-right corners where all test results are either negative or positive, respectively. The axes of true positives and false positives do not obviously translate into modifications of the pre-test probability or odds ratio, in contrast to equations (3) and (6). As motivated above, likelihood ratios provide a more useful metric when considering these factors. Therefore, examining the ROC curve after its transformation into likelihood ratio co-ordinates merits consideration.

Figure 3 shows several ROC curves before and after their conversion into likelihood ratio co-ordinates using equations (2) and (5). These ROC curves come from continuous diagnostic tests which separate normal distributions by increasing integer numbers of standard deviations. The likelihood ratio axes have been further transformed by base-ten logarithms to simplify

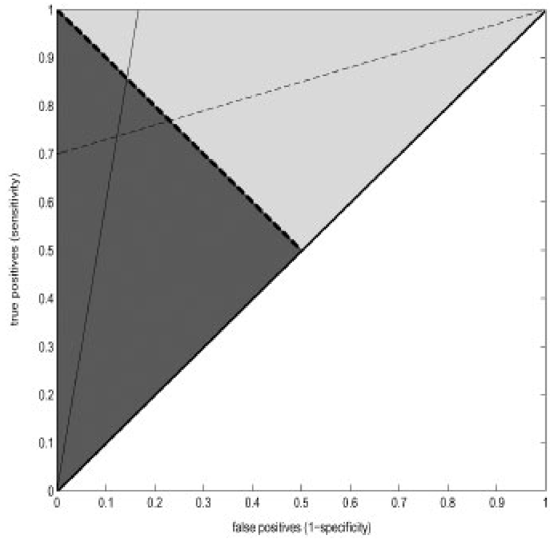

Figure 2. The potential region for ROC curves can be divided into three important regions by the lines $\mathsf{Se} = \mathsf{Sp}$ (thick dashed) and $\mathsf{Se} = 1 - \mathsf{Sp}$ (thick solid). Along these lines $|\log \mathsf{LR}_+| = |\log \mathsf{LR}_-|$. Below $\mathsf{Se} = 1 - \mathsf{Sp}$ (not shaded) $\mathsf{LR}_+ < 1$ and $\mathsf{LR}_- > 1$. Between the thick solid and dashed lines $|\log \mathsf{LR}_+| > |\log \mathsf{LR}_-|$ (darker shading). Above the thick solid and dashed lines $|\log \mathsf{LR}_-| > |\log \mathsf{LR}_+|$ (lighter shading). Points on a line through the lower-left corner (thin solid) produce the same LR_+, as given by equation (9). Points on a line through the upper-right corner (thin dashed) produce the same LR_-, as given by equation (10).

application of equation (7). In Figure 3(a) ever better ROC curves bend more towards the optimal upper-left corner, where $\mathsf{Se} = \mathsf{Sp} = 1$. Similarly, in Figure 3(b) ever better ROC curves move away from the lower-left corner, where $|\log_{10} \mathsf{LR}_+| = |\log_{10} \mathsf{LR}_-| = 0$ and towards the upper-right area where $|\log_{10} \mathsf{LR}_+|$ and $|\log_{10} \mathsf{LR}_-|$ approach infinity.

As shown in Figure 3(b) using these new co-ordinates, only the curve representing populations with means separated by three standard deviations has operating points for which both $|\log_{10} \mathsf{LR}_+| \geqslant 1$ and $|\log_{10} \mathsf{LR}_-| \geqslant 1$. Both of the other curves fall outside this area, meaning that no operating point can be chosen such that both positive and negative test results change the odds ratio by at least an order of magnitude. Such quantitative observations cannot be easily made from the curve viewed with standard axes in Figure 3(a).

Lines of constant LR_+ and LR_- provide a means of bounding all proper ROC curves which pass through their intersection. Consider an arbitrary point on an ROC curve, and plot the lines of constant LR_+ (through the lower-left corner) and constant LR_- (through the upper-right corner) which pass through that point. Such a point and its likelihood ratio lines are shown in Figure 4(a). The area between these lines (shaded in Figure 4) must contain the rest of the ROC curve [4]. Equivalently, $|\log \mathsf{LR}_+|$ must decrease monotonically while $|\log \mathsf{LR}_-|$ increases monotonically. The horizontal and vertical lines of Figure 4(b) provide an easier visual restriction than the sloped ones of Figure 4(a).

Figure 4(a) also provides a means of comparing two binary diagnostic tests [8]. As noted above, all points in the area below a line of constant LR_+ and a line of constant LR_- have

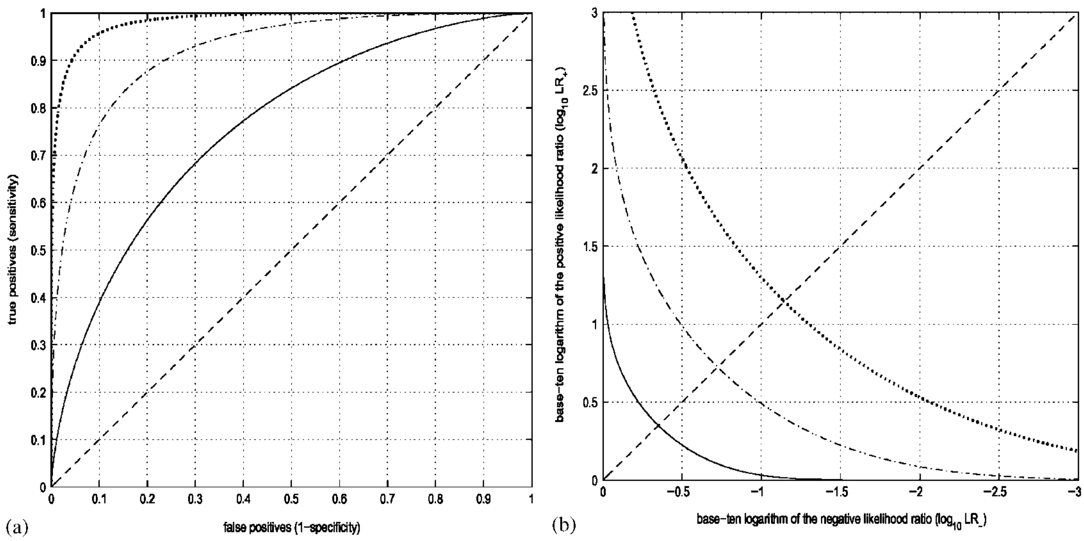

Figure 3. (a) ROC curves plotted using the traditional axes of true positives (sensitivity) and false positives (one minus specificity). The three curves correspond to normal probability density functions as in Figure 1(a), here separated by one, two, and three standard deviations (solid, dash–dot, and dotted lines, respectively). (b) ROC curves from (a) plotted using the axes of positive and negative likelihood ratios, defined by equations (2) and (5), further scaled using base-ten logarithms. The straight line (dashed) has $|\log \mathsf{LR}_+| = |\log \mathsf{LR}_-|$.

lower values of $|\log \mathsf{LR}_+|$ and $|\log \mathsf{LR}_-|$ than at the point of intersection. Likewise, all points above both lines have higher values of $|\log \mathsf{LR}_+|$ and $|\log \mathsf{LR}_-|$ than at the point of intersection. Points in the darker shaded area have a higher $|\log \mathsf{LR}_+|$ but a lower $|\log \mathsf{LR}_-|$. Points in the lighter shaded area have a lower $|\log \mathsf{LR}_+|$ but a higher $|\log \mathsf{LR}_-|$.

Therefore, to compare two binary diagnostic tests, plot one test on an ROC curve along with corresponding lines of constant LR_+ and LR_-. If the other diagnostic test lies below both lines, its $|\log \mathsf{LR}_+|$ and $|\log \mathsf{LR}_-|$ are both less than the other test's. If it lies above both lines, its $|\log \mathsf{LR}_+|$ and $|\log \mathsf{LR}_-|$ are both greater than the other test's. Otherwise, the two tests offer a tradeoff between LR_+ and LR_-. Note that Figure 4(b) eases such a comparison, as the lines become vertical and horizontal and the areas become rectangular. Such comparisons also have a relationship to regret graphs [8, 9].

The traditional ROC axes have necessary limits of 0 and 1, inclusive. The logarithm likelihood ratio axes are only bounded by 0 at one end. Choosing the second limit for each axis now becomes a matter of choice depending on the ROC curve at hand. While this means that graphs of different ROC curves might not be as comparable as with the traditional axes, upper limits of ± 3 or ± 4 (corresponding to $|\log \mathsf{LR}_+|$ and $|\log \mathsf{LR}_-|$ of 1000 or 10000, respectively) should be adequate for almost all graphs. Indeed, the ability to focus the graph on a region of interest, say a tradeoff between LR_+ and LR_- for a specific clinical situation, will make the new axes even more useful.

The AUC using traditional ROC axes ranges from $\frac{1}{2}$ for a test with no diagnostic power ($\mathsf{Se} + \mathsf{Sp} = 1$) to 1 for a 'perfect' test ($\mathsf{Se} = \mathsf{Sp} = 1$). The first case maps to the single

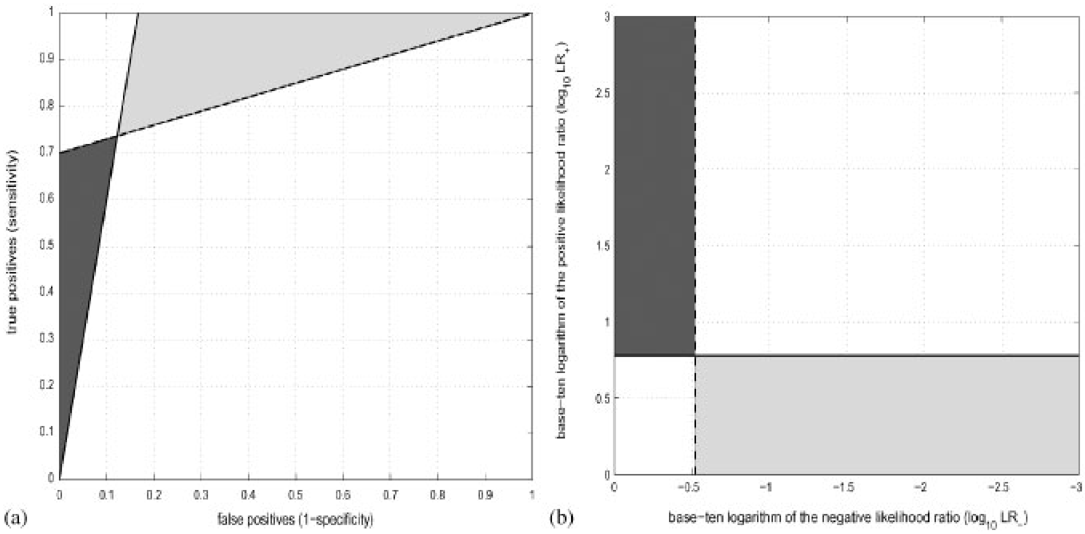

Figure 4. Proper ROC curves are bounded by lines of constant LR_+ and LR_-. Solid lines denote constant LR_+. Dashed lines denote constant LR_-. The darker shaded area covers regions with increasing $|\log \mathsf{LR}_+|$. The lighter shaded area covers regions with increasing $|\log \mathsf{LR}_-|$. (a) A proper ROC curve which passes through a point is bounded by the straight lines from $(0,0)$ and to $(1,1)$ which pass through that point. (b) A transformed proper ROC curve which passes through a point in base-ten logarithm positive and negative likelihood ratio co-ordinates is bounded by horizontal and vertical lines which pass through that point.

point $(|\log \mathsf{LR}_-|, |\log \mathsf{LR}_+|) = (0,0)$ and the second case maps to the single point $(|\log \mathsf{LR}_-|, |\log \mathsf{LR}_+|) = (\infty, \infty)$. The area under an ROC curve using transformed co-ordinates similarly ranges from 0 to infinity for these two cases. In practice, the AUC using traditional axes often lies toward the upper end of its range, where a small increase in the AUC corresponds to a vastly better diagnostic test. The AUC using the transformed axes assigns the more meaningful value of 0 to a useless diagnostic test, and does not suffer from crowding of values as tests approach 'perfection'.

CONCLUSION

Graphing the ROC curve using logarithm likelihood ratio co-ordinates offers several advantages for interpretation. First, points of no diagnostic power $(\mathsf{Se} + \mathsf{Sp} = 1)$ map to $(0,0)$. The rest of the space denotes diagnostically useful tests. Second, these co-ordinates more clearly reinforce visually that $|\log \mathsf{LR}_+|$ decreases monotonically as $|\log \mathsf{LR}_-|$ increases monotonically. Third, isocontours of constant LR_+ and LR_- fall along horizontal and vertical lines instead of sloping lines. Fourth, scaling by a base-ten logarithm allows visual quantification of the impact of a diagnostic test on equation (7), and therefore makes more apparent the tradeoffs for a given operating point of the associated binary test.

The traditional axes of true positives (sensitivity) and false positives (one minus specificity) obscure the more important measures of positive and negative likelihood ratios when examining a diagnostic test. We encourage future authors to give ROC curves using positive and negative likelihood ratio axes.

ACKNOWLEDGEMENTS

The author would like to thank Daniel T. Johnson for commenting on a draft of this paper, and David Garner and other staff at Perceptronix Medical Incorporated and the British Columbia Cancer Research Centre for stimulating thought on this topic.

REFERENCES

1. Fletcher RH, Fletcher SW, Wagner EH. *Clinical Epidemiology: The Essentials* (3rd edn). Williams & Wilkins: Baltimore, 1996; 43–74.
2. Swets JA. Measuring the accuracy of diagnostic systems. *Science* 1988; **240**(4857):1285–1293.
3. Green DM, Swets JA. *Signal Detection Theory and Psychophysics*. Wiley: New York, 1966; 7–52.
4. Egan JP. *Signal Detection Theory and ROC Analysis*. Academic Press: New York, 1975; 6–56.
5. Swets JA. *Signal Detection Theory and ROC Analysis in Psychology and Diagnostics: Collected Papers*. Lawrence Erlbaum Associates: Mahwah, NJ, 1996; 7–30.
6. Hanley JA, McNeil BJ. The meaning and use of the area under a receiver operating characteristic (ROC) curve. *Radiology* 1982; **143**(1):29–36.
7. Hilden J. The area under the ROC curve and its competitors. *Medical Decision Making* 1991; **11**(2):95–101.
8. Biggerstaff BJ. Comparing diagnostic tests: a simple graphic using likelihood ratios. *Statistics in Medicine* 2000; **19**(5):649–663.
9. Hilden J, Glasziou P. Regret graphs, diagnostic uncertainty and Youden's index. *Statistics in Medicine* 1996; **15**(10):969–986.

