

Fairness Worksheet

Question Sheet

Zak Varty

The questions on this sheet are designed to let you test your own understanding of the course content on fairness. Some questions will test basic notions, while others will encourage you to think more deeply about some of the concepts introduced this week.

A pharmaceutical company has developed three new tests to screen patients for Coeliac disease (a condition where your immune system attacks your own tissues when you eat gluten), which are rapid and cost-effective. Receiving a diagnosis of Coeliac disease usually requires multiple visits to a general practitioner or specialist. This leads to long delay times between the first appointment and diagnosis. It is hoped that this delay can be reduced by the new tests, if they are effective.

Each of the three tests were administered to n people of whom m had been diagnosed and $n - m$ had no negative reactions to consuming gluten. For each individual $i = 1, \dots, n$ and each test $j = 1, 2, 3$, let:

- Y_i be an indicator that individual i has Coeliac disease;
- $D_{j,i}$ be an indicator that test j returns a positive diagnosis for individual i ;
- A_i be a random variable denoting whether individual i belongs to subgroup “a” (children under 18), “b” (adults over 18 but under 65), or “c” (adults over 65).

Question 1: Confusion matrix terminology

Using the notation introduced previously for test 1, define and interpret in plain language:

1. The true positive rate (TPR) of the test,
2. The false positive rate (FPR) of the test,
3. The true negative rate (TNR) of the test,
4. The false negative rate (FNR) of the test.

Question 2: PPV and NPV

How do the positive predictive value and negative predictive value relate to the rates defined in Question 1?

Question 3: Calculating with Confusion Matrices

The confusion matrices for the three tests are given below.

<i>Test 1</i>	D = 1	D = 0
Y = 1	81	24
Y = 0	24	382

<i>Test 2</i>	D = 1	D = 0
Y = 1	73	55
Y = 0	32	351

<i>Test 3</i>	D = 1	D = 0
Y = 1	5	3
Y = 0	100	403

- Calculate the TPR, FPR, TNR and FNR for test 1, showing your working clearly.
- Calculate and state the TPR, FPR, TNR, and FNR for tests 2 and 3.
- Is it important that the same group of people took each of the three tests? Why or why not?
- Calculate the sensitivity and specificity for each of these tests.

Question 4: ROC curves

The receiver operating characteristic for a test plots the sensitivity of the test against 1 - specificity. This pair can be used to compare different tests for the same binary outcome. The ROC for test 1 is shown in Figure 1.

- Which region of the ROC plot corresponds to a near-optimal classifier?
- What does it mean for a classifier to have an ROC on/above/below the line $y = x$?




Figure 1: ROC plot for test 1

- c) Add the ROC for tests 2 and 3 to Figure 1.
- d) Which, if any, of the tests would you recommend that the company develops further?
- e) By altering the concentration of an enzyme in the tests, the pharmaceutical company can change the threshold at which each test gives a positive diagnosis. The ROC curve for a test interpolates ROC values for each test at a range of enzyme concentrations. The ROC curves for tests 1,2, and 3 are shown in Figure 2.

When enzyme levels are optimised for predictive performance, which test has the best results? Does this change your earlier recommendation?

- f) The company decides to further develop test 2, having considered the cost, logistics and predictive performance of all three tests. A follow-up study is used to establish the effects of age on test outcomes. Three age groups are considered: group “a” represents people under 18, group “b” represents people 18-65 years old, and group “c” those over 65.

The ROC curves for test 2 are shown for people in age-groups “a”, “b” and “c” in Figure 3.

Interpret the ROC curves shown. You should use plain language suitable for an executive summary to the directors of the pharmaceutical company.

- g) The fairness condition of error-parity ensures that the false positive rates and the false negative rates of a test are equal for all sub-groups of a protected characteristic, such as age.

Shade the area of the ROC plot in which tests satisfying the error-parity fairness condition will be located. Explain why this is the case.

- h) What does your answer to 4(g) imply about the relative predictive performance of classifiers with and without an error-parity condition?




Figure 2: ROC curves for tests 1, 2 and 3

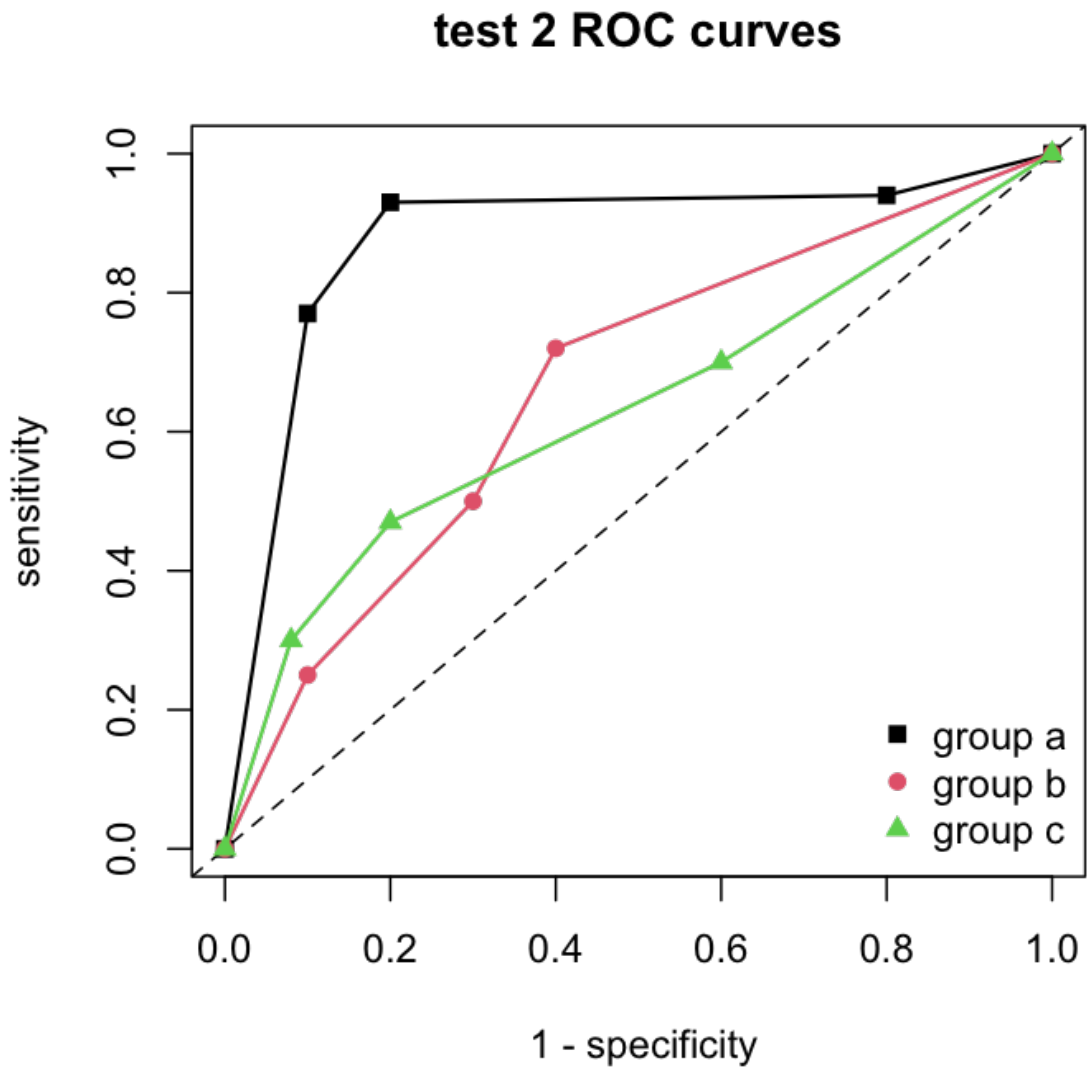


Figure 3: Sub-group specific ROC curves for test 2