

Evaluation Metrics

(find and work with a partner)


Say you google “ROC”, wanting documents about ROC curves. Say there are 100 total documents to search through, and 10 are returned to you. Of these, 7 are about ROC curves. In reality, there are 20 documents about ROC curves.

1. Complete the confusion matrix below, adding the sum of each row to the right and the sum of each column to the bottom.

		Predicted class	
		Negative	Positive
True class	Negative		
	Positive		

2. What is the precision? The recall (true positive rate)?

3. Compute the false positive rate and then plot this result on the ROC curve axes below.



4. With the same query, another search engine gives you 30 documents, 18 of which are about ROC curves. Add this point to your ROC curve above.

5. (Ethics Discussion) As a search engine user, which result would you prefer? What other factors would influence your satisfaction with the results? If this were a clinical example (test for a disease), how might your response change? What other real-world scenarios would cause you to rethink the 50% threshold?