

- Determine whether to employ the classifier or
- Optimise the classifier

Info 411, Machine Learning and Data Mining

3/39 05-Sep, 2017

05-Sep, 2017

4/39





Lecture Slides

2



Info 411, Machine Learning and Data Mining

05-Sen 2017

Lecture Slides

05-Sep, 2017

11/39

12/39

Info 411, Machine Learning and Data Mining



Scalar Characteristics Are Not Good For **Evaluating Performance**

- Scalar characteristics such as the Accuracy, F-measure, Precision and Recall do not provide enough information
- ► E.g., they don't tell you:
 - ▶ How are errors distributed across the classes?
 - ▶ How will each classifier perform in different testing conditions?
- ▶ Two numbers True Positive Rate (TPR) and False Positive Rate (FPR) are much more informative than a single number

Info 411, Machine Learning and Data Mining

And these two scalar values can make a curve!

Receiver Operating Characteristic (ROC) Curve¹

- Developed in 1950s for signal detection theory to analyse noisy signals
- ▶ Plots TPR against FPR
- Performance of each classifier represented as a point on the ROC curve:
 - changing the threshold of algorithm, sample distribution or cost matrix changes the location of the point



05-Sep, 2017

05-Sep. 2017

15/39







Lecture Slides





Info 411: Machine Learning and Data Mining, Semester 2, 2017











Lecture Slides

1.00

5

5

32/39





Lecture 8: Performance Evaluation of Classifiers

Lecture Slides

Info 411: Machine Learning and Data Mining, Semester 2, 2017





References

	C. Drummond & R. C. Holt "What ROC curves can't do (and cost curves can)" ECAI workshop on ROC analysis in artificial intelligence, 2004	
	T. Fawcett "An introduction to ROC analysis" Pattern Recognition Letters. 27 , 861–874, 2006	
	D. J. Hand & R. J. Till "A Simple Generalisation of the Area Under the ROC Curve for Multiple Class Classification Problems" Machine Learning. 45 (2): 171–186, 2001	
	D. J. Hand "Measuring classifier performance: a coherent alternative to the area under the ROC curve" Machine Learning. 77 (1): 103–123, 2009	
05-Sep, 2017	Info 411, Machine Learning and Data Mining	39 / 39