Receiver operating characteristic (ROC) curves

Statistical concepts for clinical investigators
David M. Thompson
January 14, 2014

When a clinical test or measure (in this example, CMAP duration) is observed on a continuum, an important question involves choosing a cut point or threshold that places a patient into one or the other of two diagnostic categories. The ROC curve illustrates the tradeoffs between cut points that maximize sensitivity (which can help rule out a diagnosis) and specificity (which can help rule in a diagnosis).

One can choose a cut point based on clinical criteria (for example, the desire to minimize either false positives or false negatives) or, probably less often, mathematical criteria (for example, the cut point that maximizes the sum of sensitivity and specificity).

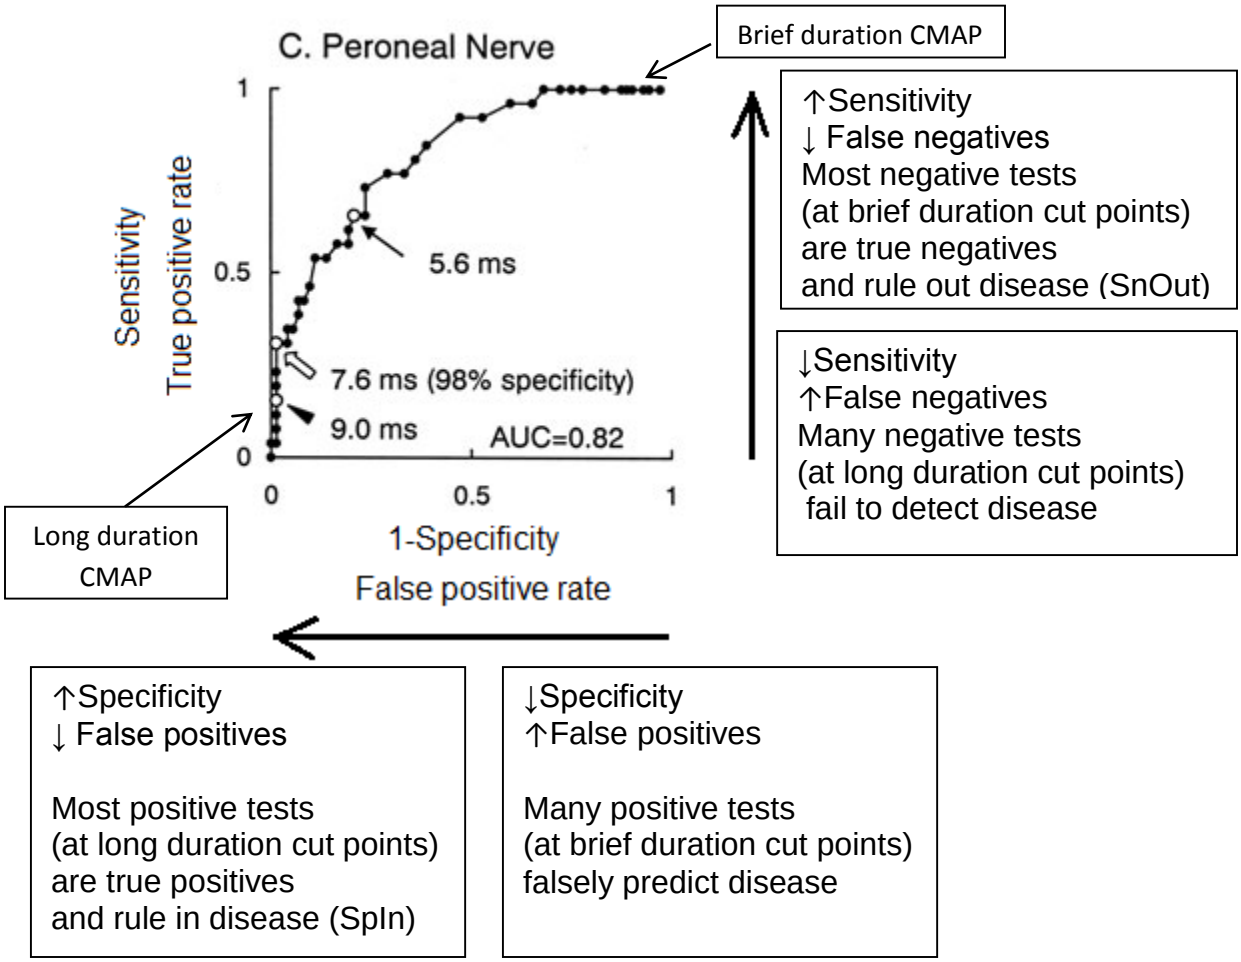

Diagram modified from Isose S, Kuwabara S, Kokubun N, Sato Y, Mori M, et al. Utility of the distal compound muscle action potential duration for diagnosis of demyelinating neuropathies. Journal of the Peripheral Nervous System 14:151–158 (2009)

The area under the ROC curve (AUC) gives a rough idea of a test's predictive power. The AUC is equivalent to the probability that, confronted with a pair of randomly chosen patients, one of whom truly has the disease of interest and the other of whom truly does not, the test will accurately identify which of the pair has the disease. The probability can approach 1 at the high end. At the low end, one hopes it is still higher than 0.5, which is equivalent to a coin toss or a random guess.