Motivation #1 – Realism

Wish to take into account

- misclassification costs
- class distribution (at testing time)

Expected cost of a classifier =

$$p \bullet miss_1 \bullet cost_1 + (1-p) \bullet miss_2 \bullet cost_2$$

p = probability that a given example is in class 1

$miss_i$ = probability the classifier will misclassify an example in class i

$cost_i$ = cost if an example in class i is misclassified

Often p and $cost_i$ are not known during learning.

Main Approach: parameterized classifier that can be adapted to current values of p and $cost_i$

Alternative: ROC Convex Hull (Provost and Fawcett) generalizes the above approach