



Pattern Recognition

Classification cost
Biometric system errors
ROC curve



Bayesian classification

- Bayes classification rule: classify x to the class w_i which has biggest posterior probability $P(w_i | x)$

$$P(w_1 | x) > P(w_2 | x) ? \quad w_1 : w_2$$

- Bayes classification rule minimizes the total probability of misclassification.

Cost of errors.

- Errors happen when samples of class 1 are incorrectly classified to belong to class 2, and samples of class 2 are classified to belong to class 1.
- The cost of making these errors can be different :

λ_1 - the cost of misclassifying samples of class 1

λ_2 - the cost of misclassifying samples of class 2



Total cost (or risk) of classification

Classification algorithm splits feature space into two decision regions:

R_1 - samples in this region are classified as being in class 1

R_2 - samples in this region are classified as being in class 2

$\int_{R_2} p(x | w_1) dx$ - the proportion of samples of class 1 being classified as class 2

$\int_{R_1} p(x | w_2) dx$ - the proportion of samples of class 2 being classified as class 1

$P(w_1) \int_{R_2} p(x | w_1) dx$ - the proportion of all input samples being class 1 but classified as being in class 2

$P(w_2) \int_{R_1} p(x | w_2) dx$ - the proportion of all input samples being class 2 but classified as being in class 1

$$\text{Cost} = \lambda_1 P(w_1) \int_{R_2} p(x | w_1) dx + \lambda_2 P(w_2) \int_{R_1} p(x | w_2) dx \quad \text{- total cost}$$



Bayesian classification

Bayesian classifier is an optimal classifier minimizing total classification cost. Such classifier is possible only if we have full knowledge about class distributions.

If $\lambda_1 P(w_1) p(x | w_1) > \lambda_2 P(w_2) p(x | w_2)$ then classify x as class 1.

If $\lambda_1 P(w_1) p(x | w_1) \leq \lambda_2 P(w_2) p(x | w_2)$ then classify x as class 2.

Alternatively, assuming non-zero terms, the class assignment is based on

testing whether $\frac{p(x | w_1)}{p(x | w_2)} > \frac{\lambda_2 P(w_2)}{\lambda_1 P(w_1)}$ or $\frac{p(x | w_1)}{p(x | w_2)} \leq \frac{\lambda_2 P(w_2)}{\lambda_1 P(w_1)}$

Decision surface $\frac{p(x | w_1)}{p(x | w_2)} = \frac{\lambda_2 P(w_2)}{\lambda_1 P(w_1)}$ separates two decision regions.

$\frac{p(x | w_1)}{p(x | w_2)}$ - likelihood ratio

$\frac{p(x | w_1)}{p(x | w_2)} > (<) \frac{\lambda_2 P(w_2)}{\lambda_1 P(w_1)}$ - likelihood ratio test



Biometric Application Types

- **Verification System (1:1)**
 - Claim is made (enrollee identity)
 - User's biometric is matched only with stored biometric of claimed enrollee
 - The decision to accept claim is made using only one matching score

- **Identification System (1:N)**
 - No claim about identity is made
 - User's biometric is matched with stored biometrics of all enrolled persons
 - The highest matching score determines the most probable enrollee
 - The decision about accepting identification attempt is made based on the matching score for that enrollee (and optionally using other matching scores too)

- **Screening**
 - Matching against a watch list
 - Opposite of verification



Errors in Verification Systems

Each verification attempt has two possibilities:

1. Genuine event - input biometrics and stored biometrics from claimed identity belong to the same person.
2. Impostor event - input biometrics is different from claimed identity biometrics.

The scores produced by matching algorithm will have distributions:

$$p_{gen}(s) = p(s \mid \text{genuine event})$$

$$p_{imp}(s) = p(s \mid \text{impostor event})$$



Errors in Verification Systems

FAR and FRR are determined by the decision rule – accept or reject results of recognition.

Usually FAR and FRR are defined using some threshold:

$$FAR(\theta) = \int_{\theta}^{\infty} p_{imp}(s) ds = P(s > \theta | \text{impostor event})$$

$$FRR(\theta) = \int_{-\infty}^{\theta} p_{gen}(s) ds = P(s < \theta | \text{genuine event})$$



Errors in Verification Systems





Figure 5.2: The non-match scores are on average lower than the match scores; in this case, the threshold T is set high to minimize False Accept.



Performance of Biometric Matchers





Estimating FAR and FRR

In contrast to estimating pdf, FAR and FRR are easily estimated:

$$FAR(t) = \int_{x>t} p(x | imp) dx \approx \frac{|\{x_i | x_i > t, x_i \text{ is impostor}\}|}{|\{x_i | x_i \text{ is impostor}\}|}$$

$$FRR(t) = \int_{x<t} p(x | gen) dx \approx \frac{|\{x_i | x_i < t, x_i \text{ is genuine}\}|}{|\{x_i | x_i \text{ is genuine}\}|}$$



ROC Curve

ROC curve connects $FAR(\theta)$ and $FRR(\theta)$ curves.

Note that they both use same θ at the same time, so we are able to construct such plot.




Figure 5.4: The ROC curve expresses the trade-off between FMR and FNMR.



Types of ROC Curve




Figure 5.5: The ROC with one probability scale in logarithmic form; on the left the FMR is expressed in logarithmic form, on the right the FNMR is in logarithmic form.




Figure 5.6: The ROC with both probability scales in logarithmic form.

Taking $\log(FAR(\theta))$
 and $\log(FRR(\theta))$
 instead of $FAR(\theta)$
 and $FRR(\theta)$ is
 reasonable if they are
 small.



Types of ROC Curve




Figure 5.7: A detection error curve with the detection rate ($1 - \text{FNMR}$), plotted against the False Alarm Rate.



Using ROC Curve




Figure 5.8: Two distinct ROC curves and a sample operating point specified with a target FMR. Matcher *b* is more accurate than *a* for all T .



Comparing ROC Curves

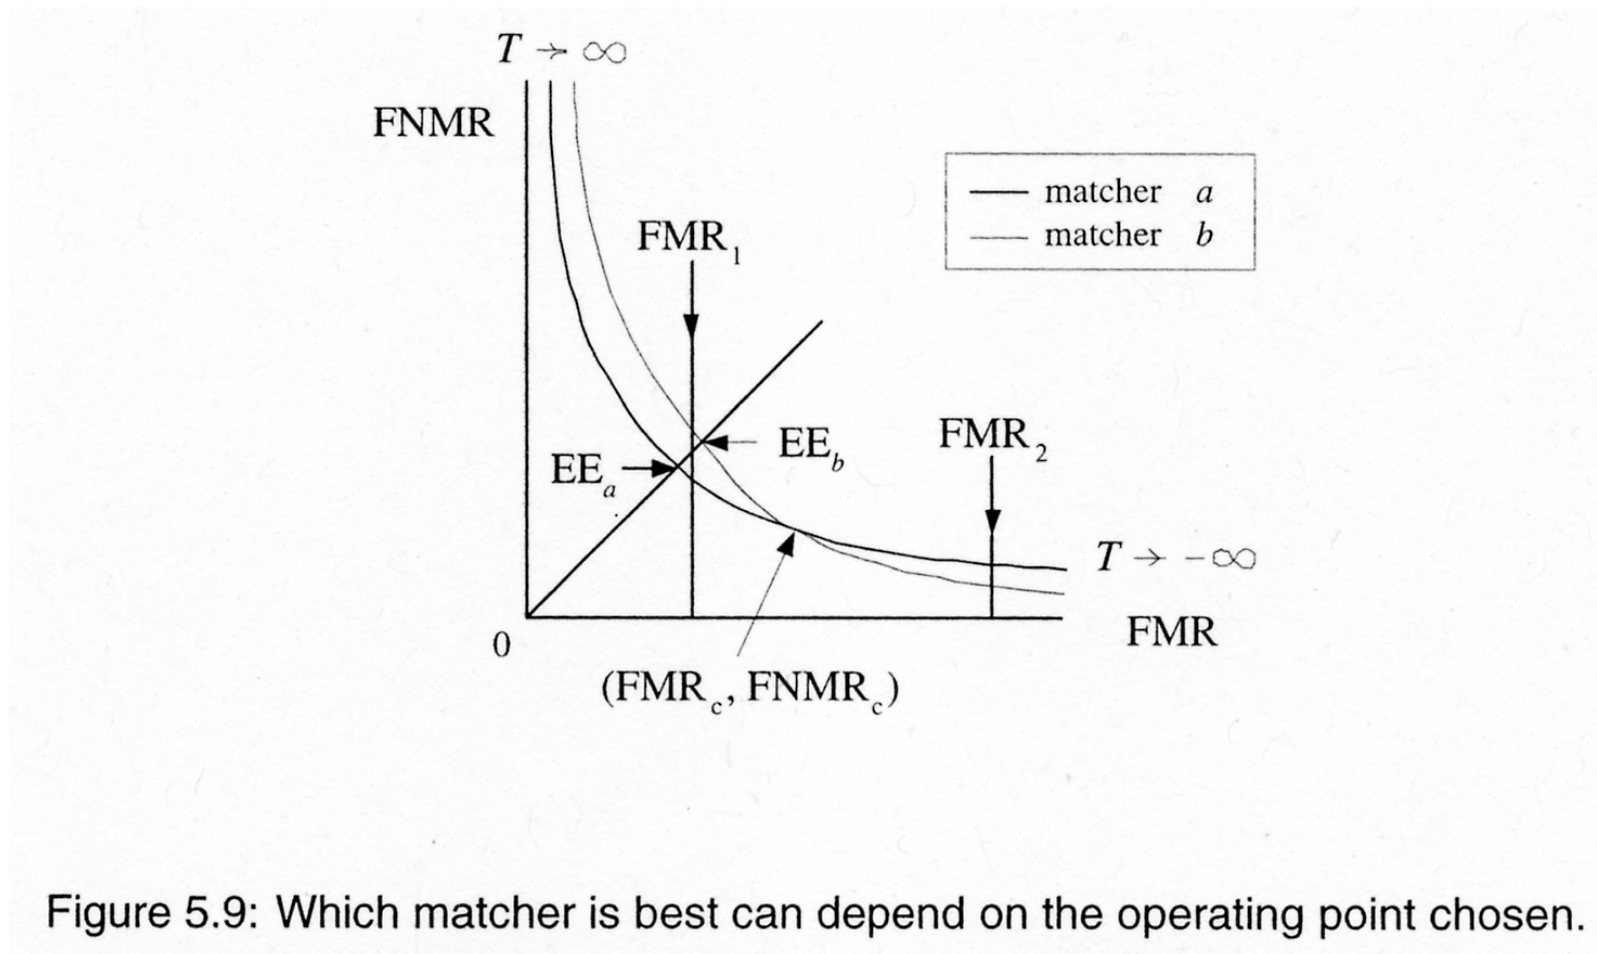


Figure 5.9: Which matcher is best can depend on the operating point chosen.



Trade-offs

Selection of the operating point in a particular application is a trade-off between security and convenience.




Figure 5.13: ROCs for Matcher *a* and *b*. Matcher *b* may be preferred for convenience and Matcher *a* for security..



Using FAR and FRR

In Bayesian framework we want to minimize total cost:

$$\begin{aligned} \text{Cost} &= C_{FA} P(\text{impostor}) P(s > \theta | \text{impostor}) \\ &\quad + C_{FR} P(\text{genuine}) P(s < \theta | \text{genuine}) \\ &= C_{FA} P_{imp} FAR(\theta) + C_{FR} P_{gen} FRR(\theta) \end{aligned}$$

Correct setting of θ in verification application requires estimating $C_1, C_2, P(\text{impostor}), P(\text{genuine})$



Example

Consider the problem of deploying biometric matcher for an amusement park admission

$C_{FA} = \$20$ - cost of accepting impostor to the park

$P_{imp} = 1\%$ - probability of impostor attempts


$C_{FR} = \$1$ - cost of rejecting genuine user

$P_{gen} = 99\%$ - probability of genuine attempts

$$Cost = C_{FA} P_{imp} FAR(\theta) + C_{FR} P_{gen} FRR(\theta)$$


$$= 20 \times .01 \times FAR(\theta) + 1 \times .99 \times FRR(\theta)$$

$$= .2 \times FAR(\theta) + .99 \times FRR(\theta)$$



Face matcher 'C' better minimizes cost

$$Cost = .2 \times FAR(\theta) + .99 \times FRR(\theta)$$




If we had more impostor attempts, say $P_{imp} = 10\%$, then matcher 'ri' would get lower cost

$$Cost = 2 \times FAR(\theta) + .9 \times FRR(\theta)$$



Comparing ROC Curves



Area under ROC curve (1-FRR vs FAR) represents the probability that random genuine score is higher than random impostor score.



Comparing ROC Curves

Compare match and non-match score densities by d-prime method:

$$d' = \frac{\mu_m - \mu_n}{\sqrt{\sigma_m^2 + \sigma_n^2}}$$




Figure 5.10: Different ROCs for two hypothetical matchers a and b with identical d' . Here Gaussian score distributions with identical means and different variances lead to the same d' but different ROCs.



Comparing ROC Curves




Figure 5.11: The minimum expected error will not generally be found at the same operating point as the Equal Error Rate.

Equal Error Rate (EER): $EER = FRR(\theta) = FAR(\theta)$
 at θ such as $FRR(\theta) = FAR(\theta)$

Minimum Total Error Rate (TER):

$$TER = \min_{\theta} FRR(\theta) + FAR(\theta)$$



Errors in Identification Systems

N people are enrolled in the database. The recognition algorithm performs N matchings with output scores:

$$s_1 > s_2 > \dots > s_N$$

(the scores are ordered by magnitude, but not by people id)

The decision algorithm usually considered:

- Accept class 1 if

$$s_1 > \theta \text{ and } \theta > s_2 > \dots > s_N$$

- Reject otherwise



Errors in Identification Systems

Other types of decisions involve selecting a subset of matched classes:

- Threshold based:

$$s_1 > s_2 > \dots > s_k > \theta$$

-select all classes bigger than threshold

- Rank –based:

-select k classes with best scores

- Hybrid:

-select based on threshold, if not successful select k classes based on rank



FNMR and FMR in Identification Systems

FNMR – False non-match rate:


$$FNMR(\theta) = FRR(\theta) = \int_{-\infty}^{\theta} p_{gen}(s) ds = P(s < \theta | \text{genuine})$$

FMR – False match rate:

$$\begin{aligned}
 FMR(\theta) &= P(\max s_i > \theta | i \text{ corresponds to all } N - 1 \text{ impostor event}) \\
 &= 1 - P(s_i < \theta | i \text{ corresponds to all } N - 1 \text{ impostor event}) \\
 &= 1 - \prod_i P(s_i < \theta | i \text{ corresponds to one impostor event}) \\
 &= 1 - \prod_i (1 - P(s_i > \theta | i \text{ corresponds to one impostor event})) \\
 &= 1 - \prod_i \left(1 - \int_{\theta}^{\infty} p_{imp}(s) ds \right) = 1 - (1 - FAR(\theta))^{N-1}
 \end{aligned}$$



FMR for different N





Errors in Identification Systems

- FMR and FNMR might not adequately describe the performance of identification systems
 - closed set / open set identification
 - rejecting all identification results might be a correct choice
 - errors are connected: impostor might be a top choice, but genuine is also higher than the threshold
- Score belonging to different classes are usually dependent, so FMR can not be effectively estimated by means of FAR
- Still no good standard for measuring identification system performance exists