

Machine learning approaches for person identification and verification

Harish Srinivasan[†], Matthew J. Beal and Sargur N. Srihari[†]

{hs32@cedar, mbeal@cse, srihari@cedar}.buffalo.edu

Department of Computer Science and Engineering,
University at Buffalo, The State University of New York, Buffalo NY, USA

[†]Center of Excellence for Document Analysis and Recognition (CEDAR), Buffalo NY

ABSTRACT

New machine learning strategies are proposed for person identification which can be used in several biometric modalities such as friction ridges, handwriting, signatures and speech. The biometric or forensic performance task answers the question of whether or not a sample belongs to a known person. Two different learning paradigms are discussed: person-independent (or general learning) and person-dependent (or person-specific learning). In the first paradigm, learning is from a general population of ensemble of pairs, each of which is labelled as being from the same person or from different persons– the learning process determines the range of variations for given persons and between different persons. In the second paradigm the identity of a person is learnt when presented with multiple known samples of that person– where the variation and similarities within a particular person are learnt. The person-specific learning strategy is seen to perform better than general learning (5% higher performance with signatures). Improvement of person-specific performance with increasing number of samples is also observed.

Keywords: machine learning, verification, identification, person-independent, person-dependent, biometrics, signature verification, speaker identification, friction ridge identification

1. INTRODUCTION

Person identification, regardless of the method, is ubiquitous in our daily lives. Two basic questions that often need to be answered are “Who are you?” (Identification) and “Are you who you claim to be?” (Verification). For example, one may have to be authenticated before gaining access to a bank, a protected site or to draw cash from an ATM. Also authentication of bank cheques,¹ questioned document examination and other biometric verifications seek a wide variety of applications. Figure 1 shows an instance where an anonymous questioned signature sample is matched against multiple known samples of a particular writer. Biometric-based person identification is considered more reliable since the biological characteristics cannot be forgotten (like passwords) and cannot be easily shared, misplaced or duplicated. The characteristics could be a person’s physical traits e.g., fingerprints, hand geometry or behavioral characteristics e.g., voice and signature.

Automatic person identification and verification can be performed using machine learning. This article focuses on the strategies for learning the identity of a person and the techniques for automatic identification/verification that can be applied with any biometric mentioned above. Two different learning techniques: *person-independent* and *person-dependent* are mentioned. The person-independent strategy learns from a large general population of biometric samples. The person-dependent strategy focuses on learning from samples of the particular person. There has been a lot of earlier work on writer and signature verification/identification²⁻⁵ and fingerprint matching and classification.⁶

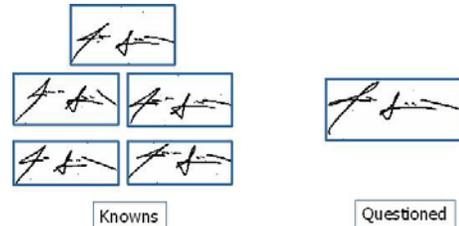


Figure 1. Signature verification with multiple knowns.

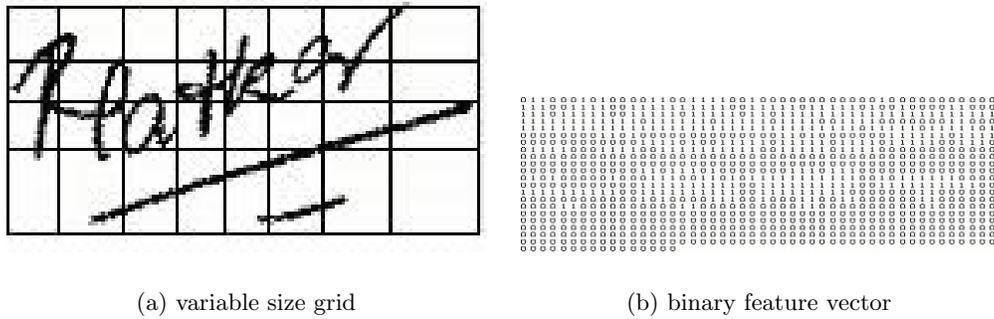


Figure 2. Feature computation signature: (a) variable size grid, and (b) binary feature vector.

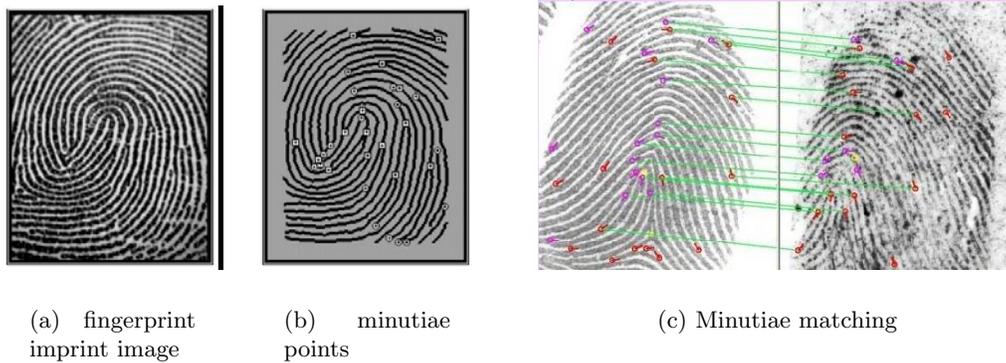


Figure 3. Feature computation and minutiae matching.

2. LEARNING STRATEGY

The paper discusses two kinds of learning strategies. The person-independent(general) learning strategy can be treated as a one-shot phase that learns from a large general population of samples. The person-dependent(person specific) learning focuses specifically on the person whose identity needs to be learnt from multiple known samples available and then answers the question if an anonymous questioned sample belongs to this person or not. These samples may be fingerprint imprint images, signature samples, speech samples etc. Both strategies operate on a set of *features* of the sample, which are described in the following section; however, we note that our machine learning paradigm is generally applicable to any biometric data with sufficiently discriminating feature elements. Before learning truly begins, these raw samples first need to be converted to feature space using a feature extractor.

2.1. Feature extraction

The *features* for a given set of samples of a particular person can be termed as a set of elements that help uniquely label the samples as belonging to that person. These elements are termed Discriminating Elements (DE's) when two samples are compared. A given person's samples can have a (possibly variable) number of DE's and the combination of the DE's have greater power in uniquely identifying the person. Features for different biometrics are different. Figure 2(a) shows an example of a signature, from which features known as Gradient, Structural and Concavity (or GSC) features have been extracted. Such features are discussed in.⁷⁻¹¹ A large number of binary features have been extracted from these, as shown in Figure 2(b), which are termed *global word shape features*.¹² Figures 3(a) and 3(b) show another example of features computed for a sample fingerprint imprint image.

The features are now considered representative characteristics of the sample. In order to compare two samples and to quantify their similarity, a similarity measure or a distance measure is used to compute a score that signifies the strength of match between the features of the two samples. This similarity measure again can be different depending upon which biometric and what features are being considered (for example, some similarity measures that can be used with binary vectors have been discussed in¹³). The similarity measure converts the data from feature space to *distance* space. Data in distance space for signatures results out of comparing two feature vectors of two signatures using a similarity measure. In a similar manner, features for fingerprints are characterized by minutiae, and Figure 3(c) shows two fingerprint imprints whose minutiae have been extracted and that are being matched. Each minutia is characterized by a triplet $\{x, y, \theta\}$, corresponding to the (x, y) coordinates of the minutia and its angular direction θ . The following procedure can be used to compute a distance measure between two fingerprint imprints. Let T be a template image with M minutiae and I be the input image with N minutiae:

$$T = \{m_1, m_2, \dots, m_M\} \quad \text{where} \quad m_i = \{x_i, y_i, \theta_i\}. \quad i = 1..M. \quad (1)$$

$$I = \{m_{1'}, m_{2'}, \dots, m_{M'}\} \quad \text{where} \quad m_{j'} = \{x_{j'}, y_{j'}, \theta_{j'}\}. \quad j = 1..N. \quad (2)$$

The distance between a minutia $m_{j'}$ in I and a minutia m_i in T can be calculated using *spatial distance* (sd) and *direction difference* (dd)

$$sd(m_{j'}, m_i) = \sqrt{(x_{j'} - x_i)^2 + (y_{j'} - y_i)^2} \quad (3)$$

$$dd(m_{j'}, m_i) = \min(|\theta_{j'} - \theta_i|, 360^\circ - |\theta_{j'} - \theta_i|) \quad (4)$$

Other methods and details for matching minutiae have been discussed elsewhere.¹⁴ Eventually, irrespective of the biometric used, one can arrive at a distance space representation of the data that characterizes the strength of match between two samples.

2.2. Person-independent learning (General learning)

As the name suggests, the person-independent strategy learns not from a specific individual but from a large general population by considering two different sets. Set *one* consists of a number of pairs of samples where each pair truly belongs to the same person. Set *two* consists of a number of pairs of samples where each pair truly belongs to different persons. Features are extracted for these samples and a similarity measure computes the distance between each pair. Let \vec{D}_S denote the vector of distances between all pairs in set *one*, which represents the distribution of distances when samples truly came from the same person. Similarly let \vec{D}_D denote the vector of distances between all pairs in set *two*, which represents the distribution of distances when samples truly came from different persons. These distributions can be modeled using known distributions such as Gaussian or Gamma. Figure 4(a) shows the typical distribution curves obtained when the distances are modeled using a Gamma distribution. The parameters that need to be learnt for such a model are typically derived from the sufficient statistics of the distribution, and are namely μ (mean) and σ (variance) for a Gaussian, or α (shape) and β (width) for a Gamma. These distributions are often referred to as same- and different-finger distributions in the field of fingerprint matching, and same- and different-writer distributions in the domain of signature or writer verification.

2.3. Person-dependent learning (Person specific learning)

When there are multiple known samples from a person, it is more intuitive to use all of that information to learn the similarities within and variation across that person's samples, and then use this information as a whole to match against any anonymous sample to test whether the sample belongs to this person or not. First, pairs of known samples are compared using a similarity measure to obtain a distribution over distances between features of samples — this represents the distribution of the variation/similarities amongst samples — for the particular person. The corresponding classification method involves comparing the questioned sample against

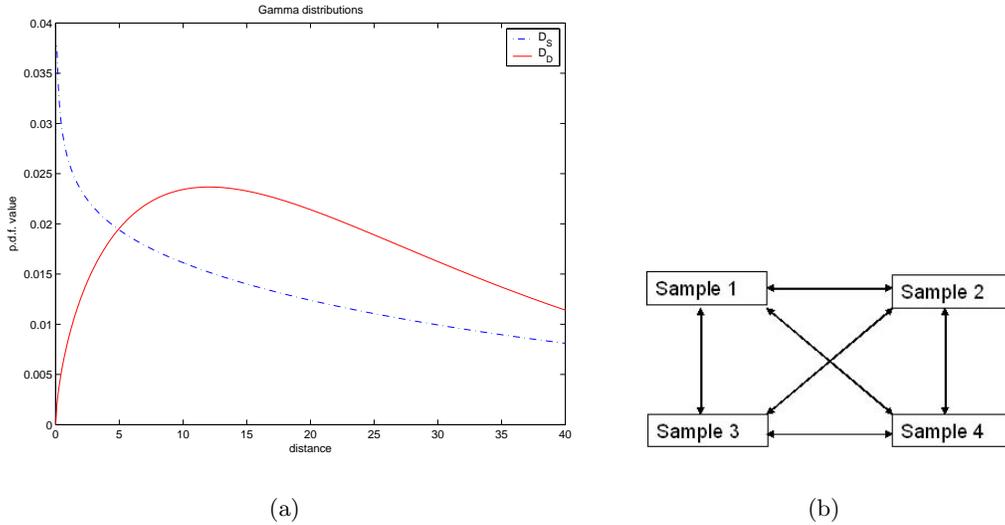


Figure 4. (a) Person-independent learning models two distributions: the distribution of distances between same persons (D_S), and the distribution of distances between different persons (D_D). In both cases we have illustrated the best fit with Gamma distributions. (b) Person-dependent learning examines the distribution of all distances between samples of that same person, as shown, to form the vector \vec{D}_W , which in this example is of length $N_{WD} = \binom{4}{2} = 6$ comparisons.

all available known samples to obtain another distribution in distance space. The Kolmogorov-Smirnov test, KL-divergence and other information theoretic methods can then be used to obtain a probability of similarity of the two distributions, which is the probability of the questioned sample belonging to the ensemble of the knowns. These methods are discussed in detail below.

2.3.1. Within-person distribution

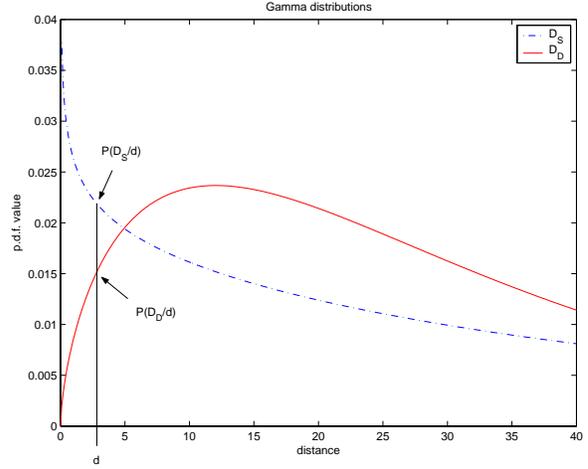
If a given person has N samples, $\binom{N}{2}$ defined as $\frac{N!}{N!(N-2)!}$ pairs of samples can be compared as shown in Figure 4(b). In each comparison, the distance between the features is computed. This calculation maps feature space to distance space. The result of all $\binom{N}{2}$ comparisons is a $\{\binom{N}{2} \times 1\}$ distance vector. This vector is the distribution in distance space for a given person. For example, in the signature verification problem this vector is the distribution in distance space for the ensemble of genuine known signatures for that writer. A key advantage of mapping from feature space to distance space is that the number of data points in the distribution is $\binom{N}{2}$ as compared to N for a distribution in feature space alone. Also the calculation of the distance between every pair of samples gives a measure of the variation in samples for that writer. In essence the distribution in distance space for a given known person captures the similarities and variation amongst the samples for that person. Let N be the total number of samples and $N_{WD} = \binom{N}{2}$ be the total number of comparisons that can be made which also equals the length of the within-person distribution vector. The within-person distribution can be written as

$$\vec{D}_W = (d_1, d_2, \dots, d_{N_{WD}})^\top \quad (5)$$

where \top denotes the transpose operation and d_j is the distance between the pair of samples taken at the j^{th} comparison, $j \in \{1, \dots, N_{WD}\}$.

3. CLASSIFICATION

The classification tasks answers the question whether or not a given questioned sample belongs to a person under consideration. Each of the above mentioned learning strategies in Sections 2.2 and 2.3 can provide their own classification answers. The person-independent method uses knowledge from a general population to answer



(a) Verification: Questioned sample verified against known sample from a person

(b) P.D.F. value of the Gamma distributions at d

Figure 5. Verification: Questioned sample verified against known sample from a person.

whether two samples belong to the same person. One may be a questioned sample and the other a known sample of that person. This task is called 1:1 verification. Section 3.1 explains the details of 1:1 verification. For the learning strategy described in Section 2.3, the corresponding classification tasks involves matching one questioned sample against multiple known samples from the person. This can be used only if multiple known samples exists that belong to that person. This type of classification can also be called person-specific classification. Section 3.2 explains the details of person-dependent or person-specific classification.

3.1. Person-independent classification

Figure 5(a) shows two fingerprint imprint images that are to be verified if they belong to the same finger. The process of such 1:1 verification starts with feature extraction and then computing the distance d between the features using a similarity measure. From the learning described in Section 2.2, the likelihood ratio defined as $\frac{P(D_S|d)}{P(D_D|d)}$ can be calculated, where $P(D_S|d)$ is the probability density function value under the D_S distribution at the distance d and $P(D_D|d)$ is the probability density function value under the D_D distribution at the distance d . If the likelihood ratio is greater than 1, then the classification answer is that the two samples do belong the same person and if the ratio is less than 1, they belong to different persons. Figure 5(b) shows how the likelihood ratio is obtained. If there are a total of N known samples from a person, then for one questioned sample N , 1:1 verifications can be performed and the likelihood ratios multiplied. In these circumstances it is convenient to deal with log likelihood-ratios rather than with just likelihood ratios. The log likelihood-ratio (LLR) is given by $\log P(D_S|d) - \log P(D_D|d)$. The decision of same-person is favored if $\log P(D_S|d) - \log P(D_D|d) > 0$, and the decision of different-person chosen if $\log P(D_S|d) - \log P(D_D|d) < 0$. When N of these 1:1 verifications are performed these LLR's are summed and then the decision is taken.

3.2. Person-dependent classification

When multiple known samples from a person are available then person-dependent learning discussed in Section 2.3 obtains the within-person distribution as explained by equation 5. A given questioned sample can now be compared against the ensemble of knowns to verify whether or not the sample belongs to the person. Figure 6 shows one questioned fingerprint sample that is to be verified when 4 known samples from a finger are available. In a similar manner a questioned signature sample can be verified if it is genuine when multiple known genuine



Figure 6. Person-dependent classification with 4 knowns.

signatures of a person are available. The classification process consists of two steps.

- (i) Obtaining Questioned Vs Known Distribution; and
- (ii) Comparison of Questioned Vs Known Distribution and Within-Person Distribution.

Questioned Vs Known Distribution In Section 2.3 and with equation 5 the within-person distribution is obtained by comparing every possible pair of samples from within the given persons samples. Analogous to this, the questioned sample can be compared with every one of the known N samples in a similar way to obtain the Questioned Vs Known Distribution. The Questioned Vs Known Distribution is given by

$$D_{QK} = (d_1, d_2, \dots, d_N)^\top, \quad (6)$$

where d_j is the distance between the questioned sample and the j^{th} known sample, $j \in \{1, \dots, N\}$.

Comparing Distributions Once the two distributions are obtained, namely the within-person distribution, denoted D_w (Section 2.3, equation 5), and the Questioned Vs Known distribution, D_{QK} (Section 3.2, equation 6), the task now is to compare the two distributions to obtain a probability of similarity. The intuition is that if the questioned sample did indeed belong to the ensemble of the knowns, then the two distributions must be the same (to within some sampling noise). There are various ways of comparing two distributions and these are described in the following sections.

Kolmogorov-Smirnov Test The Kolmogorov-Smirnov (KS) test can be applied to obtain a probability of similarity between two distributions. The KS test is applicable to unbinned distributions that are functions of a single independent variable, that is, to data sets where each data point can be associated with a single number.¹⁵ The test first obtains the cumulative distribution function of each of the two distributions to be compared, and then computes the statistic, D , which is a particularly simple measure: it is defined as the maximum value of the absolute difference between the two cumulative distribution functions. Therefore, if comparing two different cumulative distribution functions $S_{N1}(x)$ and $S_{N2}(x)$, the KS statistic D is given by $D = \max_{-\infty < x < \infty} |S_{N1}(x) - S_{N2}(x)|$. The statistic D is then mapped to a probability of similarity, P , according to equation 7

$$P_{KS} = Q_{KS} \left(\sqrt{N_e} + 0.12 + (0.11/\sqrt{N_e})D \right), \quad (7)$$

where the $Q_{KS}(\cdot)$ function is given by (see ¹⁵ for details):

$$Q_{KS}(\lambda) = 2 \sum_{j=1}^{\infty} (-1)^{j-1} e^{-2j^2 \lambda^2}, \quad \text{such that : } Q_{KS}(0) = 1, \quad Q_{KS}(\infty) = 0, \quad (8)$$

and N_e is the effective number of data points, $N_e = N_1 N_2 (N_1 + N_2)^{-1}$, where N_1 is the number of data points in the first distribution and N_2 the number in the second. The following sections discuss other methods of comparing two distributions.

Kullback-Leibler Divergence and other methods The Kullback-Leibler (KL) divergence is a measure that can be used to compare two binned distributions. The KL divergence measure between two distributions is measured in bits or nats. An information theoretic interpretation is that it represents the average number of bits that are wasted by encoding events from a distribution P with a code which is optimal for a distribution Q (i.e. using codewords of length $-\log q_i$ instead of $-\log p_i$). Jensen's inequality can be used to show that $D_{KL} = KL(P\|Q) \geq 0$ for all probability distributions P and Q , and $D_{KL} = KL(P\|Q) = 0$ iff $P = Q$. Strictly speaking, the KL measure is a *divergence* between distributions and not a distance, since it is neither symmetric nor satisfies the triangle equality). The KL divergence so obtained can be converted to represent a probability by $\exp(-\zeta D_{KL})$ (for the sake of simplicity we set $\zeta = 1$ in this article). If the divergence D_{KL} is 0, then the probability is 1 signifying that the two distributions are the same. In order to use this method and other methods discussed in the following sections it is first necessary to convert the two unbinned distributions to binned distributions with a probability associated with each bin. The KL divergence between two distributions is given in equation 9 below, where B is the total number of bins, P_b and Q_b are the probabilities of the b^{th} bin of two distributions respectively. P_{KL} denotes the probability that the two distributions are the same. Other related measures between distributions P and Q that we will examine are given in equations 10, 11 and 12

$$\text{Kullback-Leibler: } D_{KL} = \sum_{b=1}^B P_b \log\left(\frac{P_b}{Q_b}\right) \quad P_{KL} = e^{-\zeta D_{KL}} \quad (9)$$

$$\text{Reverse KL: } D_{RKL} = KL(Q\|P) = \sum_{b=1}^B Q_b \log\left(\frac{Q_b}{P_b}\right) \quad P_{RKL} = e^{-\zeta D_{RKL}} \quad (10)$$

$$\text{Symmetric KL: } D_{HKL} = \frac{1}{2}KL(P\|Q) + \frac{1}{2}KL(Q\|P) = \frac{D_{KL} + D_{RKL}}{2} \quad P_{HKL} = e^{-\zeta D_{HKL}} \quad (11)$$

$$\text{Jensen-Shannon KL: } D_{JS} = \frac{1}{2}KL\left(P\left\|\frac{P+Q}{2}\right.\right) + \frac{1}{2}KL\left(Q\left\|\frac{P+Q}{2}\right.\right) \quad P_{JS} = e^{-\zeta D_{JS}} \quad (12)$$

Combined KL and KS measure A combination of the Kolmogorov-Smirnov and Kullback-Leibler measure, denoted KLKS, has been found to outperform the individual measures as will be analyzed in the experiment section following this. The method to combine is very simple and is obtained by averaging the probabilities defined in equations 7 and 9.

$$P_{KLKS} = \frac{P_{KL} + P_{KS}}{2} \quad (13)$$

4. EXPERIMENTS AND RESULTS

4.1. Experiments

A database of off-line signatures was prepared as a test-bed.⁴ Each of 55 individuals contributed 24 signatures thereby creating 1320 genuine signatures. Some were asked to forge three other writers' signatures, eight times per subject, thus creating 1320 forgeries. One example of each of 55 genuines are shown in Figure 7. Ten examples of genuines of one subject (subject no. 21) and ten forgeries of that subject are shown in Figure 8. Each signature was scanned at 300 dpi gray-scale and binarized using a gray-scale histogram. Salt pepper noise removal and slant normalization were two steps involved in image preprocessing. The database had 24 genuines and 24 forgeries available for each writer as in Figure 8. For each test case a writer was chosen and N genuine samples of that writer's signature were used for learning. The remaining $24 - N$ genuine samples were used for



Figure 7. Genuine signature samples.

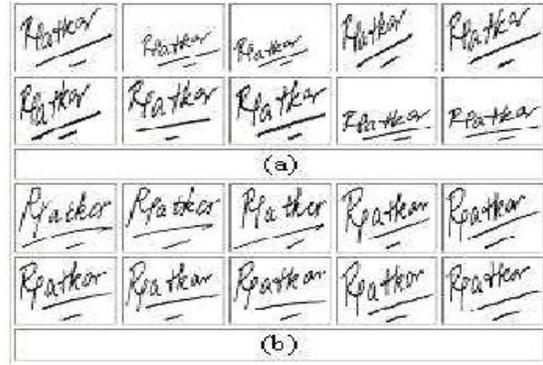


Figure 8. Samples for one writer: (a) genuines and (b) forgeries.

testing. Also 24 forged signatures of this writer were used for testing. Two different error types can be defined for any biometric person identification problem. False reject rate (Type 1) is the fraction of samples classified as not belonging to the person when truly they were from that person. False acceptance rate (Type 2) is the fraction of samples classified as belonging to the person when truly the samples were not from that person. In the domain of signatures, Type 1 is the fraction of samples classified as forgeries when truly they were genuine and Type 2 the fraction of samples classified as genuine when truly they were forgeries.

4.2. Person-independent(General) method

The classification decision boundary discussed in Section 3.1 is given by the sign of the log likelihood-ratio, LLR, $\log P(D_S|d) - \log P(D_D|d)$. A modified decision boundary can be constructed using a threshold α , such that $\log P(D_S|d) - \log P(D_D|d) > \alpha$. When α is varied, we can plot ROC curves as shown in Figure 9. The different subplots in the figure correspond to the ROC curves as the number of known samples is increased from 5 to 20. For each plot, the total error rate defined as $(\text{False acceptance} + \text{False reject})/2$ is minimum at a particular value of α . This is the best setting of α for the specified number of known samples, denoted the *operating point*, and is indicated with an asterisk '*'. When 20 samples are used for learning the error rate is approximately 79%. Figure 10 shows the distribution of LLR's when the questioned samples were genuine and when they were forgeries. A larger region of overlap indicates a higher error rate.

4.3. Person-dependent method

The person-dependent classification discussed in Section 3.1 mentioned six different statistics for comparing the two distributions to obtain a probability of match between the questioned sample and the ensemble of knowns. In order to measure error rates for this classification technique, once again a decision needs to be made based on the probability of whether or not the questioned sample belongs to the ensemble of knowns. If the probability of match $> \alpha$, then the decision is in favour of the questioned signature to be genuine, and if the probability of match $< \alpha$, the decision is in favor of a forgery (this α should not be confused with that used in the person-independent method). By varying the parameter α , once again ROC curves (False Accept vs. False Reject) can be plotted for each of the six measures. The best setting of α is termed as the *operating point*. This setting of α corresponds to the least total error rate possible. Note that the ROC curves are plotted for the test data set and the operating point determined on them. These test data set can be considered as a validation set that helps to determine the operating point. In the curve, the operating point is the point closest to the origin. Table 1 shows the least total error rate possible when different number of known samples were used for training for each of the 6 different measures. Figure 12(a) shows the same table as a graph comparing the different measures and it can be seen that the combined KL and KS measure performs the best. The reason for this can be intuitively explained

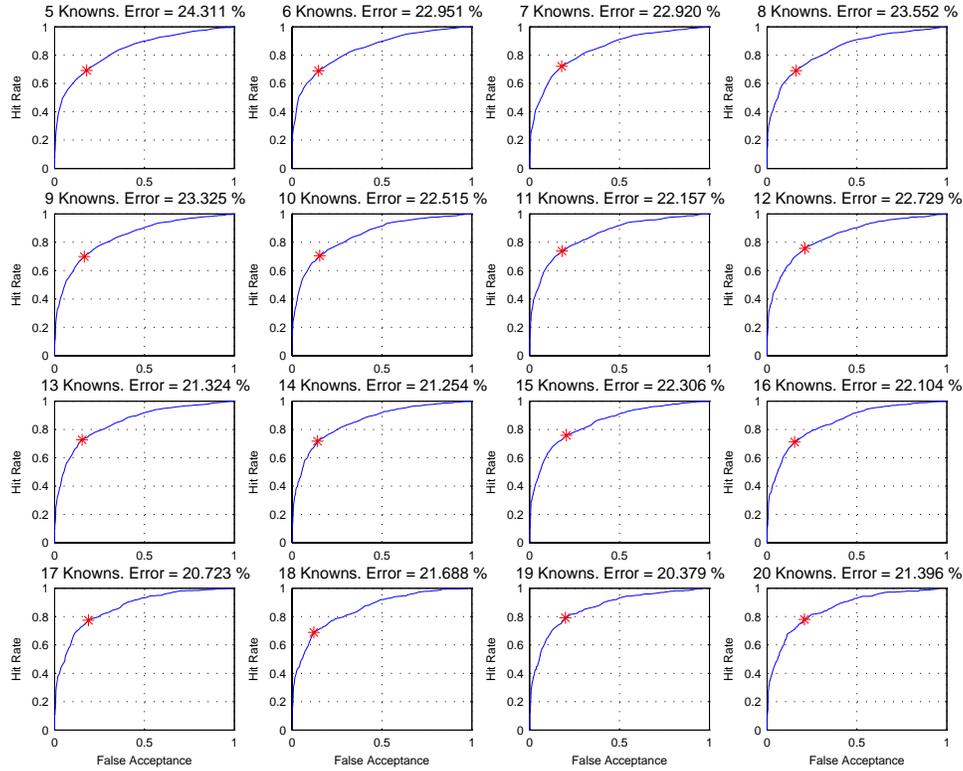


Figure 9. ROC curves parameterized by α is varied. Each subplot is titled with the number of knowns used for training and the optimum error rate that is possible. The asterisk ‘*’ denotes the optimal operating point α for that model.

by examining Figure 11(a), where it can be seen that the KS statistic has low false accept rates whereas the *KL* statistic has low false reject rates. The combination of these two in the KL and KS measure works the best.

Figure 12(b) shows how the operating point (best setting of α) varies with the number of known samples used. It can be seen that in order to obtain the least total error rate, the value of α changes with the number of knowns for certain measures. The value of α explains a lot about what each statistic learns from the known samples. For example, the high value of α for the KS statistic when large numbers of known samples were used explains that the KS statistic focuses on learning the variation amongst the known samples. Presence of large known samples accounts for greater variation amongst them. Hence if KS focuses on learning the variation, then almost every questioned sample ends up receiving a high probability of match as the majority of questioned samples (genuines and forgeries) invariably fall within the variation. Thus by setting a high value of α the decision that a sample is truly genuine is made only if probability is really high. In simple terms this means that when more samples are used for training the KS statistic will declare a sample as genuine only if the probability of match is really high. In contrast to this measure the KL measure captures the similarities amongst the known samples a lot. This is evident by the low value of α for large number of knowns. Presence of large number of samples accounts for observing more similarities. The KL measure focuses on learning the similarities amongst the samples and it returns a high probability of match very rarely and only when every similarity that is learnt amongst the known samples is present in the questioned sample. Hence the majority of questioned sample receive a low probability of match by the KL measure. To counter this a low value of α ensures that the KL measure will declare a sample as forgery only if the probability of match is really low. Similar comments can be made about other measures and it is important to note that those measures for which the operating point does not vary with the number of knowns and those which are around the 50% mark can be a useful property. This basically shows that

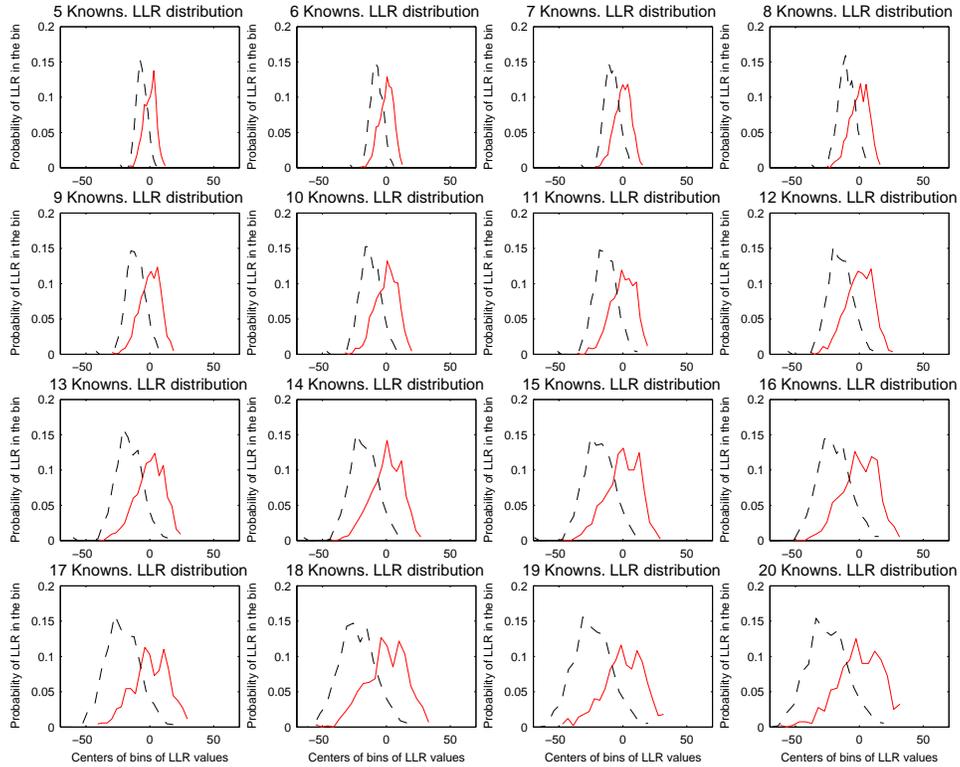
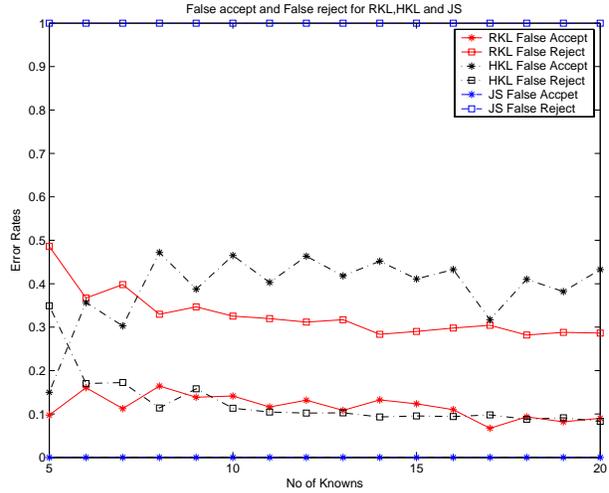
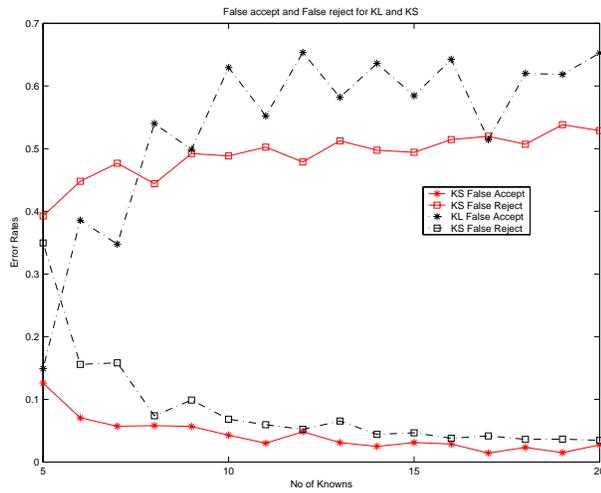


Figure 10. LLR's obtained from each test case plotted as histograms. The probability (y-axis) that the LLR falls into a range of LLR values (x-axis) is shown for the results of truly genuine (solid) and forgery cases (dotted). Each subplot corresponds to training on a different number of knowns.

irrespective of the number of knowns used for training, one can make a decision using the same operating point, and also if the operating point is around the 50% mark there is an equal range of probabilities across which the two different decisions fall. And it is also intuitive that the combined KL KS measure has this fine property. It can be seen that the operating point for the combined KL and KS measure is closest to the 50% mark amongst other measure and is also independent of the number of known samples to some extent. Proceeding with the conclusion that the combined KL KS measure has a few desired properties and also outperforms other measures in terms total error rate, we can now consider allowing rejections to reduce the error rates even further. Consider probabilities between $.5 - \beta$ and $.5 + \beta$ for some $\beta > 0$ as the region for reject probabilities. No decision is made if $.5 - \beta < Probability < .5 + \beta$. This can significantly reduce the total error rate. Figure 13 shows the total error rate as it the rejection percentage is changed by changing the value of β . This analysis enables the operator to select a value of β that will induce a certain rejection rate and in turn result in a certain desired error rate. For example, in order to obtain a error rate of 10% with 20 knowns in this data set one should set β to .15 and that accounts for 35% reject rate. Similarly for an error rate of 5% for 20 knowns, β needs be set to .30 which accounts for 62% reject rate.

5. CONCLUSIONS

We have discussed two different machine learning approaches for person identification/verification. Our learning strategies involve mapping from feature space to distance space, which makes the strategy independent of the biometric under consideration provided there is a similarity measure available to compute a distance between two features. An exhaustive set of experiments were done for the problem of automatic signature verification



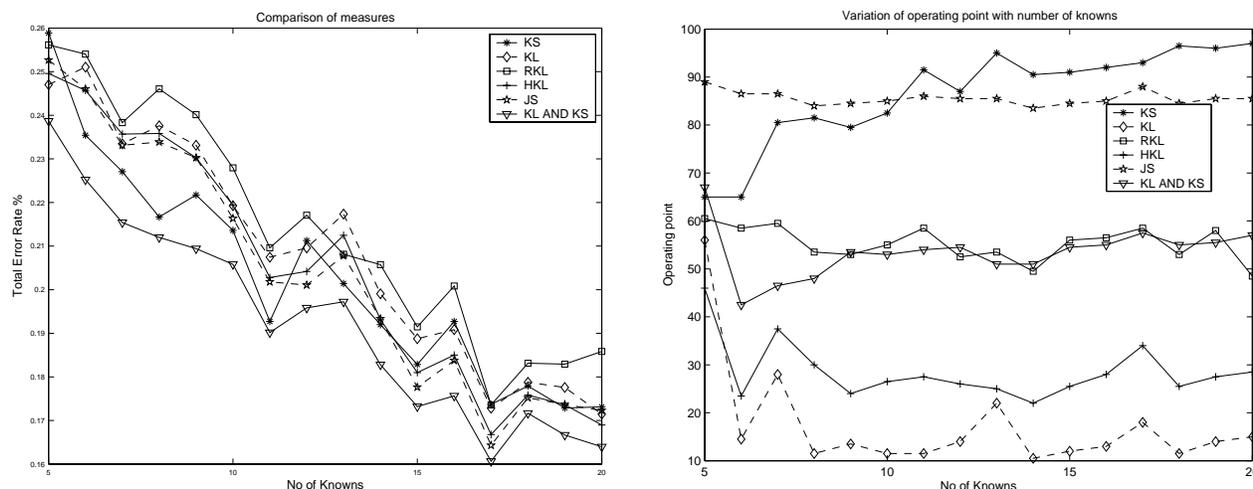
(a) False acceptance and false reject rates are plotted for the KL and the KS statistic as the number of knowns used for training is increased.

(b) False acceptance and false reject rates are plotted as in (a), for the other measures used: RKL, HKL, and JS. The operating point is fixed at 50%.

Figure 11.

No. of Knowns	KS	KL	RKL	HKL	JS	KL and KS
5	25.88	24.70	25.61	24.96	25.26	23.87
6	23.54	25.10	25.40	24.57	24.60	22.52
7	22.71	23.35	23.83	23.57	23.31	21.54
8	21.67	23.76	24.60	23.58	23.39	21.20
9	22.17	23.31	24.01	23.03	23.03	20.94
10	21.36	21.93	22.79	21.94	21.63	20.58
11	19.27	20.74	20.96	20.28	20.18	19.02
12	21.13	20.96	21.71	20.42	20.10	19.58
13	20.14	21.73	20.81	21.25	20.78	19.72
14	19.06	20.03	20.84	19.46	19.33	18.41
15	18.28	18.88	19.15	18.10	17.76	17.32
16	19.27	19.08	20.08	18.50	18.38	17.56
17	17.37	17.28	17.36	16.68	16.43	16.07
18	17.79	17.88	18.31	17.58	17.52	17.17
19	17.39	18.09	18.42	17.75	17.37	16.97
20	17.31	17.15	18.58	16.90	17.23	16.40

Table 1. Error rates for signature verification. Measures are Kolmogorov-Smirnov (KS), Kullback-Leibler (KL), reverse KL (RKL), symmetrized KL (HKL), Jensen-Shannon (JS), and combined KL and KS (KL and KS). These are graphed in Figure 12(a).



(a) Comparison of measures. The total error rate is plotted for the different measures as the number of knowns used for training is increased. The combined KL and KS measure (lowest trace) outperforms other measures (see text).

(b) Variation of operating point of the different measures, as a function of the number of knowns used for training. Operating point corresponds to the best setting of the decision boundary (probability) for classifying samples as genuine or forgery.

Figure 12.

method of scanned documents and results presented. As the number of samples used for training is increased, the error rate decreases and the person dependent (person specific) approach performs better than the person independent (general) approach by approximately 5% in total error rate. The domain of signatures and the related experiments performed is just one example of how the learning strategy and the classification techniques can be used. It suffices to say that the same can be applied to other biometrics also.

Acknowledgments: The authors wish to thank Chen Huang, Vivek Shah, and Pavitra Babu for their assistance. This work was supported in part by the National Institute of Justice grant 2004-IJ-CX-K030

REFERENCES

1. R. Plamondon and S. N. Srihari, "On-line and offline handwriting recognition: A comprehensive survey," *IEEE Transactions on Pattern Recognition and Machine Intelligence* **22(1)**, pp. 63–84, 2000.
2. R. Plamondon and G. Lorette, "On-line and offline handwriting recognition: A comprehensive survey," *IEEE Transactions on Pattern Recognition and Machine Intelligence* **22(1)**, pp. 63–84, 2000.
3. S. N. Srihari, S. Cha, H. Arora, and S. Lee, "Individuality of handwriting," *Journal Of Forensic Sciences*, pp. 856–872, 2002.
4. M. K. Kalera, B. Zhang, and S. N. Srihari, "Off-line signature verification and identification using distance statistics," *Proc. of the Int. Graphonomics society Conference*, pp. 228–232, 2003.
5. S. N. Srihari, A. Xu, and M. K. Kalera, "Learning strategies and classification methods for off-line signature verification," *Proc. of the 7th Int. Workshop on Frontiers in handwriting recognition(IWHR)*, pp. 161–166, 2004.
6. S. Pankanti, S. Prabhakar, and A. K. Jain, "On the individuality of fingerprints," *Computer Vision and Pattern Recognition(CVPR'01)*, 2001.
7. B. Fang, C. H. Leung, Y. Y. Tang, K. W. Tse, P. C. K. Kwok, and Y. K. Wong, "Off-line signature verification by the tracking of feature and stroke positions," *Pattern recognition* **36**, pp. 91–101, 2003.

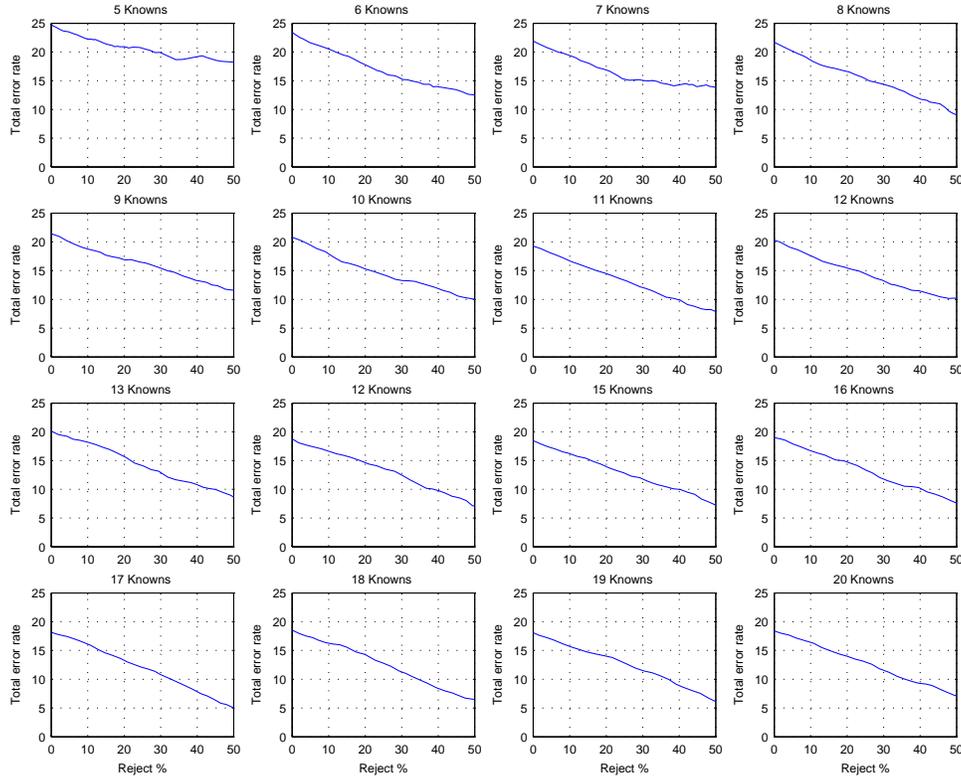


Figure 13. Error rates as the percentage of allowed rejected (no decision) cases increases. The rejection rate is indirectly controlled by varying the β which assigns the probability region $50 - \beta$ and $50 + \beta$ where no decisions are made and considered as rejects. The different subplots show the plots for different number of knowns used for learning. We have plotted only the trend for the combined KL and KS measure.

8. S. Lee and J. C. Pan, "Off-line tracking and representation of signatures," *IEEE Transactions on Systems, Man and Cybernetics* **22**, pp. 755–771, 1992.
9. R. Sabourin and R. Plamondon, "Preprocessing of handwritten signatures from image gradient analysis," *Proc. of the 8th Int. Conference on Pattern Recognition*, pp. 576–579, 1986.
10. M. Ammar, Y. Yoshido, and T. Fukumura, "A new effective approach of off-line verification of signatures by using pressure features," *Proc. of the 8th Int. Conference on Pattern Recognition*, pp. 566–569, 1986.
11. G. Srikanth, S. Lam, and S. Srihari, "Gradient based contour encoding for character recognition," *Pattern Recognition* **7**, pp. 1147–1160, 1996.
12. B. Zhang, S. N. Srihari, and C. Huang, "Word image retrieval using binary features," in *SPIE*, E. H. B. Smith, J. Hu, and J. Allan, eds., *Document Recognition and Retrieval XI* **5296**, pp. 45–53, 2004.
13. B. Zhang and S. Srihari, "Properties of binary vector dissimilarity measures," *Cary, North Carolina*, September, 2003.
14. D. Maltoni, D. Maio, A. Jain, and S. Prabhakar, *Handbook of Fingerprint Recognition*, Springer, New York
15. W. H. Press, B. P. Flannery, S. A. Teukolsky, and W. T. Vetterling, *Numerical Recipes in C: The Art of Scientific Computing*, Cambridge University Press, 1992.