GRIP-PATTERN RECOGNITION: APPLIED TO A SMART GUN

Xiaoxin Shang
Grip-pattern recognition: applied to a smart gun
Ph. D. Thesis University of Twente, December 2008
ISBN: 978-90-365-2732-3

Typeset by the author with LaTeX

The work described in this thesis was performed at the Signals and Systems group of the Faculty of Electrical Engineering, Mathematics and Computer Science, University of Twente, The Netherlands. It was part of the Secure Grip research project funded by Technology Foundation STW under project number TIT.6323.

GRIP-PATTERN RECOGNITION: APPLIED TO A SMART GUN

DISSERTATION

to obtain
the degree of doctor at the University of Twente,
on the authority of the rector magnificus,
prof. dr. W.H.M. Zijm,
on account of the decision of the graduation committee,
to be publicly defended
on Friday 19 December 2008 at 15:00

by

Xiaoxin Shang

born on 15 February 1979
in Xi'an, China.

This dissertation has been approved by:
the promotor: prof. dr. ir. C.H. Slump
the assistant promotor: dr. ir. R.N.J. Veldhuis

Contents

Chapter 1

Introduction

Abstract. *This chapter provides a general introduction to this thesis. First, the background of this research will be presented. Next, we will discuss motivation of the research, basic working principle and requirements of the system, our research questions, comparison to the work done by others, and previous work done for this research. Then, the context of this research, the Secure Grip project, will be briefly described. Finally, the research approaches and outline of this thesis will be presented.*

1.1 Smart guns in general

The operation of guns by others than the rightful users may pose a severe safety problem. In particular, casualties occur among the police officers, whose guns are taken during a struggle and used against themselves. Research in the United States, for example, has shown that approximately 8% of police officers killed in the line of duty, were shot with their own guns [1]. One of the solutions to this problem is the application of a smart gun. "Smart gun" is a phrase used throughout this thesis for the concept of weapons that have some level of user authorization capability. Today, there are a number of smart guns under research and available on the market. According to the technologies used in the recognition system, the smart guns can be categorized into three types - lockable guns, self-locking guns,

and personalized guns. They will be described briefly below.

Lockable guns

Most of the lockable guns rely on a simple trigger lock and only the rightful user has the key to activate it [2]. This ensures that the others, who do not have access to the key, are not able to fire the gun if it is deactivated. See Figure 1.1(a) for an example of a lockable gun. It is recommended that while on-duty the officer should carry an activated weapon, because otherwise there is high probability of not being able to handle the key during a life or death engagement. Therefore, the primary value of these systems is the safe storage of a weapon in an off-duty environment, for example, the home of the officer or at the police station [3].

An improvement has been made by some firearm companies that utilize a radio frequency (RF) transmitter to activate and deactivate the weapon [3]. This makes it possible for an officer to deactivate his gun if it is taken away from him. However, use of the gun can still not be secured, especially if the rightful user is disabled prior to being disarmed. For example, an officer activates the weapon at the beginning of the shift, and thus anyone getting control of the weapon can fire it unless the officer deactivates the weapon. In this scenario it is highly likely that during a struggle and attempt to retain the weapon, the officer will not be able to deactivate the weapon, and may even forget to attempt deactivation [3]. Besides, communication between the RF transmitter and the gun could be vulnerable to radio jamming or interference.

Self-locking guns

Compared to lockable guns, the self-locking guns are more secure because of their self-locking feature. These weapons are able to fire, only when the internal locking mechanism is released by placing a token close to a spot on the weapon [4], [5]. There are currently several examples of self-locking guns on the market. Figure 1.1(b), for example, shows a self-locking gun that can only be fired when a magnetic ring is held close to activate it. The ring is usually worn on a finger of the police officer. The gun is enabled and ready to fire if the police officer holds it; while it is deactivated if it has

(a)

(b)

Figure 1.1: (a) A lockable gun with a secured trigger by lock. (b) A self-locking gun with a magnetically secured trigger.

been taken away from the police officer. Compared to a lockable gun a self-locking gun is more advanced, particularly because no additional action is needed from the police officer to deactivate it. However, since these systems are controlled by a magnetic or electronic key, they are in general not personalized. Therefore, they may be vulnerable to illegal access with

forged or stolen keys [3]. Besides, when such a token is used the officer may be tempted to keep a distance to the attacker when the gun has been taken away, thus not be able to operate effectively.

Another type of self-locking gun utilizes the technology of Radio Frequency Identification (RFID), by implanting a RFID tag in one's arm, for example. In this way the weapon is personalized. However, since the system operates on radio frequencies, communication between the RFID tag and the gun could be disrupted by radio jamming or interference [6].

Personalized guns

An attractive solution to making the recognition system of a smart gun personalized, is biometric recognition. Biometrics measures and analyzes the physiological or behavioral characteristics of a person for identification and verification purposes. It associates an individual with a previously determined identity based on who one is or what one does [7]. Since many physiological or behavioral characteristics are distinctive to each person, the biometric identifiers based on these characteristics are personal. Also, many biometric recognition system have the advantage, that only minimal or even no additional action of the user is required. That is, in many applications the biometric recognition can be transparent. Here, the transparency contributes to not only convenience but also safety of the user, since explicit actions may be forgotten in stressful situations. This approach has been taken up by a small number of parties, both industrial and academic, who proposed a number of solutions.

The most well known biometric technology is probably fingerprint recognition. A fingerprint consists of a series of ridges and valleys on the surface of a finger tip. The uniqueness of a fingerprint can be determined by the pattern of ridges and valleys as well as the minutia points [8]. Personal recognition based on fingerprint has been investigated for centuries and the validity of this technology has been well established. Fingerprint recognition can be applied for the design of a smart gun, as described in [9], for example. Also, the biometric recognition in this application can be made transparent. A person implicitly claims that he or she is authorized to use a gun by holding it and putting a finger on a sensor for recognition. When the gun is released it is deactivated automatically. However, finger-

print recognition requires a clear scanned image of the fingerprint, and this is not practical in some cases because a gun can be used in all types of weather and various situations. If the user has dirty hands or is wearing gloves, for example, fingerprint recognition will be impossible. Also, the position of the fingerprint sensor must be personalized, otherwise the user would feel uncomfortable when holding the gun, due to the hard surface of the sensor. Additionally, since ambidextrous use of a gun should be allowed, it would be required that the sensors be installed on both sides of the grip of the gun, which would lead to an uncomfortable holding of the gun in any case.

Another example of a smart gun, based on a biometric recognition system, is the one based on voice recognition [10]. Such a recognition system has the following drawbacks. First, it is not reliable in noisy environments. Second, it is not practical in many situations, where a police officer is required to perform a task without being noticed. This will certainly increase risk to the life of the police officer. Third, since a user needs to speak to activate the gun, the recognition is not transparent. Besides, one may even forget to speak in a stressful situation.

1.2 Smart gun based on grip-pattern recognition

Motivation of the research

In this thesis we propose and analyze a biometric recognition system as part of a smart gun. The biometric features used in this system are extracted from a two-dimensional pattern of the pressure, exerted on the grip of a gun by the hand of a person holding it. This pressure pattern will be referred to as the grip pattern. We chose to use the grip patterns for the recognition system of a smart gun mainly for the following reasons. First, the grip-pattern recognition in our system can be made transparent, as in the case of fingerprint recognition. By holding the gun, the user implicitly claims that he or she is authorized to fire it. The biometric data are also presented implicitly, when the grip is settled. Second, we expected that a trained user, such as a police officer will have a grip pattern that is constant in time when holding the gun. This expectation was based on both discussions with experts, and the results from investigations done by

the New Jersey Institute of Technology. Figure 1.2, for example, shows that an experienced user holds the gun with the fingers always placed in the same positions [11]. Third, hand geometry is a factor that contributes to the grip pattern. This is an accepted biometric modality, which performs reasonably well. The feasibility of hand geometry recognition based on the contour of a hand has been demonstrated, for example, in [12], [13], [14], [15] and [16].

Also, a smart gun based on grip-pattern recognition does not have the drawbacks of a smart gun based on fingerprint or voice recognition. First, compared to the one based on fingerprint recognition, the performance of our system based on grip-pattern recognition is less sensitive to dirt and weather conditions, and the system can be made to work when a person wears gloves. In addition, the user will not feel uncomfortable when holding the gun with either hand. This is because the pressure sensor can be fully embedded in the grip of the gun, where it is appropriately protected against wear and tear. Second, compared to voice recognition the grip-pattern recognition in our system is transparent, as described above, and the performance of our system will not be affected by a noisy environment.

Basic working principle and requirements of the recognition system

Figure 1.3(a) shows the prototype of the smart gun. The grip of the gun is covered with a sensor sheet, capable of measuring the static pressure pattern as a function of position when the gun is being held. The sensor, used for measuring the hand-grip patterns, was a 44 by 44 piezo-resistive pressure sensor made by Tekscan Inc. [17]. Figure 1.3(b) shows an example of a grip-pattern image used in our system. Note that in a final version the sensor sheet will have to be appropriately protected against wear and tear.

During an enrollment phase a template, which is a representation of the grip pattern of the rightful user, is securely stored in the gun. As in the case of the recognition system using a fingerprint, by holding the gun a person implicitly claims that he or she is authorized to fire it. Then, the grip pattern of this person is measured and compared to the template stored in the gun. Only if they are sufficiently similar may the gun be fired by this person. Otherwise, it remains deactivated. This ensures that only

the rightful user may fire the gun, and not someone else who takes the gun away from the rightful user.

In our work, the type of biometric recognition is verification. Verification is a type of one-to-one comparison, in which the biometric system attempts to verify the identity of an individual by comparing a new biometric sample with an earlier stored template. If the two samples match, the biometric system confirms that the person is who he or she claims to be. The probability of falsely accepted patterns of impostors is called the false-acceptance rate. Its value is one if all impostor patterns are falsely accepted; and zero if none of the impostor patterns is accepted. The probability of falsely rejected patterns of the genuine users is called the false-rejection rate. The verification performance is quantified by the false-acceptance rate and the false-rejection rate.

To our grip-pattern recognition system, an important requirement is a very low false-rejection rate, rendering it highly unlikely that the rightful user is not able to fire the gun. As one can imagine it would be unacceptable if a police officer would not be able to fire his or her own gun. Currently, in The Netherlands the official requirement is that the probability of failure of a police gun be lower than 10^{-4}. Therefore, in our work the false-rejection rate for verification must remain below this value. Under this precondition, the false-acceptance rate should be minimized. We think that the acceptable false-acceptance rate should be lower than 20%. Please note that this requirement is fairly special since most biometric systems, in contrast, aim for a certain value of the false-acceptance rate and minimize the false-rejection rate.

Another requirement is that the recognition system has to be able to cope with all variations in grip patterns that may occur in practice. First, modern weapons are ambidextrous and, therefore, both the right-handed and left-handed use of the gun should be allowed. Also, wearing gloves should not hamper the operation. One possible solution to these problems might be to store multiple templates for a single user. In this case the recognition procedure is actually identification, where a one-to-many comparison is performed. Finally, the system should be robust to different grip patterns in stressful situations. Therefore, it should also be tested extensively in realistic situations. However, as preliminary research of this application and also due to the large number of topics to investigate, our

work mainly focuses on the feasibility of the grip pattern as a biometric and the development of a prototype verification system. Particularly, no research has been done so far on performance of the system in stressful situations, mainly because such realistic situations are difficult to create in an experimental setting.

Research questions

Through our research we would like to answer the following questions.

- Whether and to what extent the grip patterns are useful for identity verification?
- Whether the grip patterns of police officers are, as expected, more stable than those of untrained subjects?
- Whether grip-pattern recognition can be used for a police gun? Here we set the acceptable false-acceptance rate, at the false-rejection rate at 10^{-4}, to be lower than 20%.

Comparison to other work

Grip-pattern recognition has been investigated by the New Jersey Institute of Technology [18], [19], by Belgian weapon manufacturer FN Herstal, and by us [20], [21], [22], [23], [24], [25], [26], [27], [28]. The only results reported on this topic, besides ours and those in the patent [18], were published in [19], which method differs from the one reported by us in various aspects. First, in [19] the dynamics of the grip-pattern prior to firing are used, while in our approach recognition is based on a static grip pattern at the moment when one is just ready to fire. Second, in [19] only 16 pressure sensors are used: one on the trigger, 15 on the grip of the gun. These sensors are piezo-electric sensors, producing 16 time signals. We apply a much larger resistive sensor array, which produces a pressure image. Third, the recognition methods of both systems differ. In [19] a method based on neural networks analyzing the time signals is presented, which seems to be trained for identification, whereas we apply likelihood-ratio based verification [29], [30].

Another difference is the way that both systems have been evaluated. In [19] the data were collected from 4 shooters, while we used data from 39 trained police officers and 27 untrained users. The recognition results are, unfortunately, difficult to compare because in [19] the recognition rates obtained in an identification experiment were presented, while we present the equal-error rates in a verification experiment, which are more relevant for the final application. The equal-error rate is the value of the false-acceptance rate, when the verification system is tuned in such a way that the false-acceptance rate and the false-rejection rate are equal.

Previous work

A first, preliminary version of a grip-pattern recognition system was described in [31] and [21], in terms of its design, implementation and evaluation. An initial collection of grip patterns was gathered from a group of mostly untrained subjects, with no experience in shooting. The experimental results indicate that the hand-grip patterns contain useful information for identity verification. However, this experiment has limitations. First, all the experimental results were based on grip-pattern data collected in only one session. That is, there was no time lapse between the data for training and testing for each subject. Therefore, it brings little insight into the verification performance in a more realistic situation, where there is usually a time interval between data enrollment and recognition. Second, no data were collected from the police officers, who are the target users of the smart gun. Thirdly, the performance target of a minimized false-acceptance rate at the false-rejection rate equal to 10^{-4} was not reached.

1.3 Context of this research: the Secure Grip project

Our work is part of the Secure Grip project, where the main research question addressed is whether the hand pressure exerted while holding an object can be used to reliably authenticate or identify a person. This project is sponsored by the Technology Foundation STW, applied science division of

NWO and the technology programme of the Ministry of Economic Affairs. It consists of work by three parties described as follows.

- Recognition algorithm development. Mainly, this includes development of the recognition algorithms and collection of the grip-pattern data that can be used for design, validation, and optimization of the recognition algorithms. This work has done in Signals and Systems group of Department of Electrical Engineering at University of Twente, the Netherlands. It is the main content of this thesis.

- Security architecture design for biometric authentication systems. This research consists of work in three directions. First, the feasibility of cryptography based on noisy data has been investigated. The use of noisy data as key material in security protocols has often been suggested to avoid long passwords or keys. Second, threat modelling for biometric authentication systems has been done, which is an essential step for requirements modelling of the systems. Third, a practical solution to the problem of secure device association has been proposed, where biometric features are used to establish a common key between the pairing devices. The research has been done in Distributed and Embedded Security group of Department of Computer Science at University of Twente, the Netherlands.

- Sensor development. Mainly, this includes design and implementation of the pressure sensor, which is customized to the grip of the gun. Twente Solid State Technology (TSST), a Dutch sensor manufacturer with an interest in a market for pressure sensors, is responsible for this part of work.

1.4 Research approaches and outline of the thesis

This thesis proposes and analyzes a biometric recognition system as part of a smart gun. The remainder of this thesis is organized as follows. Chapter 2 describes the procedure of grip-pattern data collection. Meanwhile the purpose for each session of collection will be explained.

The heart of our proposed recognition system is a likelihood-ratio classifier (LRC). There were mainly two reasons for making this choice. First, the likelihood-ratio classifier is optimal in the Neyman-Pearson sense, i.e., the false-acceptance rate is minimal at a given false-rejection rate or vice versa, if the data have a known probability density function [29], [30]. Since our task is to minimize the false-acceptance rate of the system at the false-rejection rate equal to 10^{-4}, the likelihood-ratio classifier will be well-suited for this requirement. Second, experiments for grip-pattern recognition were done earlier with data collected from a group of subjects who were untrained for shooting. The verification results were compared using a number of classifiers, respectively. It was shown that the verification results based on the likelihood-ratio classifier were much better than those based on all the others [31]. Initial verification results based on a likelihood-ratio classifier will be presented and analyzed in Chapter 3, using data collected from the police officers. A major observation of this chapter is that the verification performance of the grip-pattern recognition system degrades strongly, if the data for training and testing have been recorded in different sessions with a time lapse. Further analysis will show that this is mainly attributed to the variations of pressure distribution and hand position between the probe image and the gallery image of a subject. That is, it turned out to be different from our expectation that the grip patterns of trained users are constant in time, and in reality data drift occurs in classification. However, it was also observed that the hand *shape* remains constant for the same subject across sessions.

Based on the characteristics of grip-pattern images, the verification performance may be improved by modelling the data variations during training of the classifier, or reducing the data variations across sessions, or extracting information of the hand shapes from images. As one of the solutions, a technique called double-trained model (DTM) was applied. Specifically, we combined the data of two out of three collection sessions for training and used the data of the remaining session for testing, so that the across-session data variations were better modelled during training of the classifier.

Next, in Chapter 4, two methods will be described to improve the performance of across-session verification. First, to reduce the data variation caused by the hand shift, we applied template-matching registration (TMR) as preprocessing prior to classification. Second, for comparison, the

maximum-matching-score registration (MMSR) was also implemented. In MMSR a measured image is aligned to such a position, that the matching score between this image and its template image is maximized. It was found that TMR is able to effectively improve the across-session verification results; while MMSR is not. However, the hand shift measured by MMSR was proved particularly useful in discriminating imposters from the genuine users. This inspired the application of a fused classifier, based on both the grip pattern and hand shift obtained after registration.

A novel approach, Local Absolute Binary Patterns (LABP), as image preprocessing prior to classification will be proposed in Chapter 5. With respect to a certain pixel in an image, the LABP processing quantifies how its neighboring pixels fluctuate. This technique can not only reduce the across-session variation of the pressure distribution in the images, but also it is capable of extracting information of the hand shape from an image.

Note that the grip-pattern verification in previous chapters are all based on the mean-template comparison (MTC), where a test image is compared to the mean value of all the training samples of a subject as the template. In Chapter 6, the verification results based on this method will be compared to those based on another method for comparison, the maximum-pairwise comparison (MPWC), where a test image is compared to each training sample of a subject and the sample which is the most similar to it is selected as the template. Particularly, these two methods will be compared in terms of resulting in a lower false-acceptance rate at the required false-rejection rate of 10^{-4} of the system.

During data collection sometimes a couple of lines of pixels in the grip-pattern images were found missing. It was caused by some damage in the cable of the prototype. Since in practice there can be various factors causing missing lines in a grip-pattern image while using a smart gun, such as damages to the hardware part, restoration of missing lines in the grip patterns is meaningful and necessary. In Chapter 7 we present an approach to restoration of the missing lines in a grip-pattern image, based on null-space error minimization.

The grip-pattern verification in previous chapters are all based on a likelihood-ratio classifier and low-level features, i.e. the pressure values measured as the output of the sensor. Next, we investigate the grip-pattern verification based on a number of alternative approaches. First, in Chap-

ter 8 experimental results will be presented and analyzed, based on high-level features extracted from the grip-pattern images, i.e. the physical characteristics of the hand or hand pressure of a subject. Not like the low-level features which are based on the raw data, the high-level features are based on an interpretation. Figure 1.4 illustrates the high-level features with an example. Second, in Chapter 9 verification performance of the system using the Support Vector Machine (SVM) classifier will be evaluated [32], [33]. The support vector machine classifier has been proved more capable of coping with the problem of data drifting, than other pattern-recognition classifiers [34], [35], [36], [37], [38].

All the experiments for grip-pattern verification described in previous chapters are done such, that the grip patterns used for training and those used for enrolment come from the same group of subjects. In Chapter 10 the verification performance of grip-pattern recognition is investigated in a realistic situation, where the data for training and for enrolment come from different groups of subjects.

Finally, conclusions of work presented in this thesis will be drawn in Chapter 11.

For easy comparison of the verification performances of different combinations of classifiers and algorithms, Figure 1.5 shows the verification results in different cases, using the grip patterns recorded from the police officers. 'FAR*ref*' and 'EER' represent the false-acceptance rate at the false-rejection rate equal to 10^{-4} and the equal-error rate, respectively.

(a) Markers

(b) Trained user

(c) Whole population

Figure 1.2: (a) Markers on fingertips of a hand holding a gun. (b) Scatter plots of finger markers for a trained user. (c) Scatter plots of finger markers for the whole set of subjects.

(a)

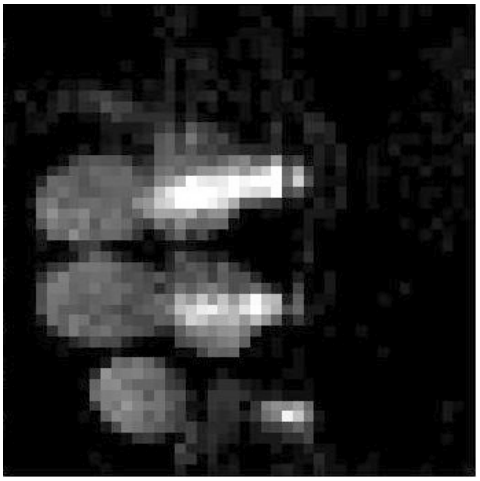

(b)

Figure 1.3: (a) Prototype of the smart gun. (b) An example of grip-pattern image.

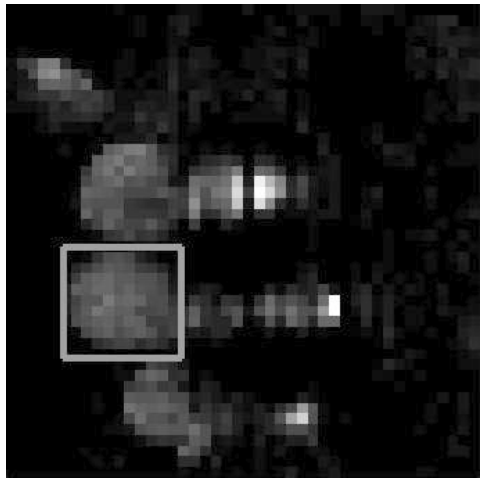

Figure 1.4: The marked rectangular subarea shows the tip section of ring finger on a grip-pattern image.

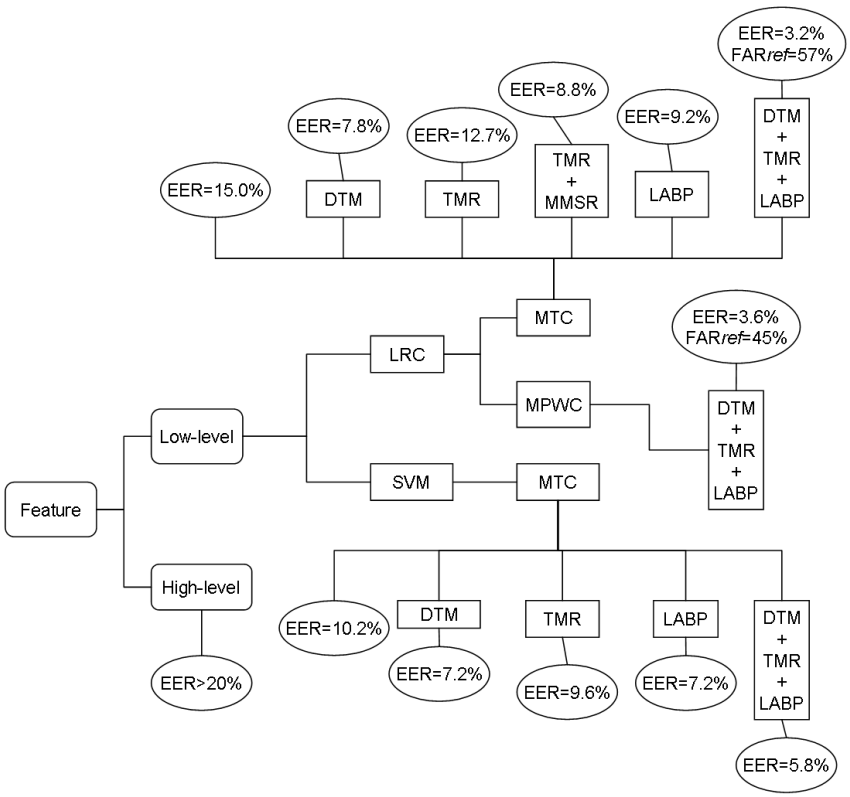

Figure 1.5: Verification results using different classifiers and algorithms.

Chapter 2

Grip-pattern data collection

Abstract. *This chapter describes the procedure of grip-pattern data collection. Furthermore, the purpose of each session of collection will be explained.*

2.1 Data collection in the first stage

We collected the grip-pattern data in two stages of the research. In the first stage, the grip-pattern images were recorded from a group of police officers in three sessions, with approximately one and four months in between. In total, 39 subjects participated in both the first and second sessions, with 25 grip-pattern images recorded from each of them. In the third session, however, the grip-pattern images were collected from 22 subjects out of the same group, and each subject contributed 50 images. In each session, a subject was asked to pick up the gun, aim it at a target, hold it, say "ready" as a signal for the operator to record the grip-pattern image, and then release the gun after the recording was finished. For each subject, this procedure was repeated till all the samples of him or her were recorded. With these data we investigated the characteristic of grip-pattern images and, based on it, developed the verification system. The experiments of this stage of research were done such, that the grip-pattern data for training the classifier were the same as those for verification enrollment.

Table 2.1: Summary of collection sessions of grip-pattern data.

Source	Name
Police	Session 1, 2, 3
Untrained	Session 4, 5, 6

2.2 Data collection in the second stage

In order to investigate the verification performance of grip-pattern recognition in a more realistic situation, where the data for training the classifier and for verification enrollment came from different groups of subjects, we needed to collect more data. In the second stage of research, therefore, the grip-pattern images were recorded from a group of people, who work or study in the Signals and Systems group of Department of Electrical Engineering at University of Twente, the Netherlands. These people were untrained subjects with no experience in shooting. The reason that we did not collect grip patterns again from the police officers was twofold. First, the sessions of data collection from the police officers turned out to be rather hard to arrange, due to administrative difficulties. Second, the across-session variations of grip patterns collected from untrained subjects proved similar, to those from the police officers. There were three collection sessions in a row, with a time lapse of about one week in between. In every session, each of the same group of 27 subjects contributed about 30 grip-pattern images. With all the grip-pattern data, collected from both the police officers and untrained subjects, the verification performance of grip-pattern recognition was investigated.

According to their order of occurrence, in this thesis we name the three sessions of grip-pattern data collected from the police officers: "Session 1", "Session 2" and "Session 3", respectively. And those recorded from the untrained people of the University of Twente are named: "Session 4", "Session 5" and "Session 6", respectively. See Table 2.1 for a summary of the collection sessions. The experiments described in Chapter 3 up to Chapter 9 are based on the data collected from the police officers. In Chapter 10 the experiments are done using the grip patterns, recorded in all six sessions.

Chapter 3

Grip-pattern verification by likelihood-ratio classifier[1]

Abstract. *In this chapter initial verification results using a likelihood-ratio classifier are presented and analyzed, with the data collected from a group of police officers. A major observation is that the verification performance degrades strongly, if the data for training and testing have been recorded in different sessions with a time lapse. This is due to the large variations of pressure distribution and hand position, between the training and test images of a subject. That is, it turned out to be different from our expectation that the grip patterns of trained users are constant in time, and in reality data drift occurs in classification. In addition, since the likelihood-ratio classifier is probability density based, it is not robust to the problem of data drift. However, it has also been observed that the hand shape of a subject remains constant, in both the training and test images. Based on these analyses, solutions have been proposed to improve the verification performance.*

[1]This chapter is based on [26].

3.1 Introduction

The main question in our research is whether or not a grip-pattern image comes from a certain subject. Verification has been done using a likelihood-ratio classifier, which is also the heart of our proposed recognition system. Initially, the reason for making this choice was twofold. First, the likelihood-ratio classifier is optimal in the Neyman-Pearson sense, i.e., the resulting false-acceptance rate is minimal at a given false-rejection rate or vice versa, if the data have a known probability density function [29], [30]. Since our research goal is to minimize the false-acceptance rate for verification at the false-rejection rate equal to 10^{-4}, the likelihood-ratio classifier is well-suited for this task. Second, experiments for grip-pattern recognition were done earlier with data collected from a group of subjects who were untrained for shooting, as described in [31]. The verification performance was compared using a number of classifiers. It was shown that the verification performance based on the likelihood-ratio classifier was much better than that based on the other types of classifiers [31].

This chapter presents and analyzes the performance of grip-pattern verification, using the likelihood-ratio classifier. Based on the analysis, solutions will be proposed to improve the verification performance. The remainder of this chapter is organized as follows. Section 3.2 describes the verification algorithm based on the likelihood-ratio classifier. The experimental results will be presented and discussed in Section 3.3. Finally, conclusions will be given in Section 3.4.

3.2 Verification algorithm based on a likelihood-ratio classifier[2]

We arrange the pixel values of a measured grip-pattern image into a (in this case $44 \times 44 = 1936$-dimensional) column vector $\mathbf{z}$. The likelihood ratio $L(\mathbf{z})$ is given by

$$L(\mathbf{z}) = \frac{p(\mathbf{z}|\mathrm{c})}{p(\mathbf{z}|\bar{\mathrm{c}})}, \tag{3.1}$$

[2]This section is based on [21].

where $p(\mathbf{z}|\bar{\mathrm{c}})$ is the probability of $\mathbf{z}$, given that $\mathbf{z}$ is not a member of class c. Since we assume an infinite number of classes, the exclusion of a single class does not change the distribution of the feature vector $\mathbf{z}$. That is, the distribution of $\mathbf{z}$, given that $\mathbf{z}$ is not a member of c, equals the prior distribution of $\mathbf{z}$ as

$$p(\mathbf{z}|\bar{\mathrm{c}}) = p(\mathbf{z}). \tag{3.2}$$

As a result, the likelihood ratio given in (3.1) can be expressed as

$$L(\mathbf{z}) = \frac{p(\mathbf{z}|\mathrm{c})}{p(\mathbf{z})}. \tag{3.3}$$

We further assume that the grip-pattern data are Gaussian. The class c is thus characterized by its local mean vector μ_c and local covariance matrix $\mathbf{\Sigma}_\mathrm{c}$; while the total data are characterized by the total mean vector μ_T and total covariance matrix $\mathbf{\Sigma}_\mathrm{T}$. The subscripts 'c' and 'T' represent 'class c' and 'total', respectively. Rather than the likelihood ratio, we use a matching score $\mathbf{z}$ derived from the log-likelihood ratio, under the assumption that the grip-pattern data are Gaussian. It is given by

$$\begin{aligned} M(\mathbf{z}) \;=\; & -\frac{1}{2}(\mathbf{z}-\mu_\mathrm{c})'\mathbf{\Sigma}_\mathrm{c}^{-1}(\mathbf{z}-\mu_\mathrm{c}) \\ & + \frac{1}{2}(\mathbf{z}-\mu_\mathrm{T})'\mathbf{\Sigma}_\mathrm{T}^{-1}(\mathbf{z}-\mu_\mathrm{T}) \\ & - \frac{1}{2}\log|\mathbf{\Sigma}_\mathrm{c}| + \frac{1}{2}\log|\mathbf{\Sigma}_\mathrm{T}|, \end{aligned} \tag{3.4}$$

where $'$ denotes vector or matrix transposition. If $M(\mathbf{z})$ is above a preset threshold, the measurement is accepted as being from the class c. Otherwise it is rejected. That is, the threshold determines the false-rejection rate and false-acceptance rate for verification.

In practice the mean vectors and covariance matrices are unknown, and have to be estimated from the training data. The number of training samples from each class should be much greater than 1936, the number of elements in a feature vector in our case. Otherwise, the classifier would become overtrained, and the estimates of $\mathbf{\Sigma}_\mathrm{T}$ and $\mathbf{\Sigma}_\mathrm{c}$ would be inaccurate. However, we cannot acquire this large number of measurements, as training

of the classifier would become very impractical. Moreover, even if enough measurements could be recorded, the evaluation of (3.4) would, with 1936-dimensional feature vectors, still be too high a computational burden.

These problems are solved by whitening the feature space and at the same time reducing its dimensionality, prior to classification. The first step is a Principal Component Analysis (PCA) [39], determining the most important dimensions (with the greatest variances) of the total data. The principal components are obtained by doing a singular value decomposition (SVD) on the matrix $\mathbf{X}$, the columns of which are the feature vectors taken from N_{user} subjects in the training set. The data matrix $\mathbf{X}$ has $N_{\text{raw}} = 1936$ rows and N_{ex} columns. Now let us assume that $\mathbf{X}$ has zero column mean. If necessary, the column mean has to be subtracted from the data matrix prior to the SVD. As a result of the SVD the data matrix $\mathbf{X}$ is written as

$$\mathbf{X} = \mathbf{U}_{\mathbf{X}}\mathbf{S}_{\mathbf{X}}\mathbf{V}_{\mathbf{X}}^{'}, \tag{3.5}$$

with $\mathbf{U}_{\mathbf{X}}$ an $N_{\text{raw}} \times N_{\text{ex}}$ orthonormal matrix spanning the column space of $\mathbf{X}$, $\mathbf{S}_{\mathbf{X}}$ an $N_{\text{ex}} \times N_{\text{ex}}$ diagonal matrix of which the (non-negative) diagonal elements are the singular values of $\mathbf{X}$ in descending order, and $\mathbf{V}_{\mathbf{X}}$ an $N_{\text{ex}} \times N_{\text{ex}}$ orthonormal matrix spanning the row space of $\mathbf{X}$. The whitening and the first dimension-reduction step are achieved as follows. Let the $N_{\text{raw}} \times N_{\text{PCA}}$ matrix $\mathbf{U}_{\text{PCA}}$ be the submatrix of $\mathbf{U}$ consisting of the first $N_{\text{PCA}} < N_{\text{ex}}$ columns. Furthermore, let the $N_{\text{PCA}} \times N_{\text{PCA}}$ matrix $\mathbf{S}_{\text{PCA}}$ be the first principal $N_{\text{PCA}} \times N_{\text{PCA}}$ submatrix of $\mathbf{S}$. Finally, let the $N_{\text{ex}} \times N_{\text{PCA}}$ matrix $\mathbf{V}_{\text{PCA}}$ be the submatrix of $\mathbf{V}$ consisting of the first N_{PCA} columns. The whitened data matrix with reduced dimensions is now given by

$$\mathbf{Y} = \sqrt{N_{\text{ex}} - 1}\mathbf{V}_{\text{PCA}}^{'}. \tag{3.6}$$

The resulting dimension N_{PCA} must be chosen such that only the relevant dimensions, i.e. with sufficiently high corresponding singular values are kept. A minimum requirement is that all diagonal element of $\mathbf{S}_{\text{PCA}}$ are strictly positive. The corresponding whitening transform is

$$\mathbf{F}_{\text{white}} = \sqrt{N_{\text{ex}} - 1}\mathbf{S}_{\text{PCA}}^{-1}\mathbf{U}_{\text{PCA}}^{'}. \tag{3.7}$$

The total column mean of $\mathbf{Y}$ is zero, and the total covariance matrix of $\mathbf{Y}$ is an identity matrix $\mathbf{I}_{\text{PCA}} = \frac{1}{N_{\text{ex}}-1}\mathbf{Y}\mathbf{Y}^{'}$.

The whitened matrix $\mathbf{Y}$ can now be used to estimate the within-class covariance matrices. Here we make a simplifying assumption that the within-class variations of all classes are characterized by one within-class covariance matrix. The reason is that often not enough data from each class are available to reliably estimate individual within-class covariance matrixes. First, the subjects' contributions to the training data can be ordered such that

$$\mathbf{Y} = (\mathbf{Y}_1, \ldots, \mathbf{Y}_{N_{\text{user}}}), \tag{3.8}$$

with $\mathbf{Y}_i$ the whitened data from class i. The column mean ν_i from $\mathbf{Y}_i$ estimates the mean feature vector of class i after whitening. The matrix

$$\mathbf{R} = (\mathbf{Y}_1 - \nu_1, \ldots, \mathbf{Y}_{N_{\text{user}}} - \nu_{N_{\text{user}}}) \tag{3.9}$$

contains all variations around the means.

We now proceed to estimate a diagonalized version of the within-class covariance matrix after whitening. A second SVD on $\mathbf{R}$ results in

$$\mathbf{R} = \mathbf{U_R} \mathbf{S_R} \mathbf{V_R}^{'}, \tag{3.10}$$

with $\mathbf{U_R}$ an $N_{\text{PCA}} \times N_{\text{PCA}}$ orthonormal matrix spanning the column space of $\mathbf{R}$, $\mathbf{S_R}$ an $N_{\text{PCA}} \times N_{\text{PCA}}$ diagonal matrix of which the (non-negative) diagonal elements are the singular values of $\mathbf{R}$ in descending order, and $\mathbf{V_R}$ an $N_{\text{ex}} \times N_{\text{PCA}}$ orthonormal matrix spanning the row space of $\mathbf{R}$. The within-class covariance matrix can be diagonalized by pre-multiplying $\mathbf{R}$ by $\mathbf{U}_{\mathbf{R}}^{'}$. For the resulting, diagonal, within-class covariance matrix, further denoted by $\mathbf{\Lambda_R}$, we have

$$\mathbf{\Lambda_R} = \frac{1}{N_{\text{ex}} - 1} \mathbf{S}_{\mathbf{R}}^2. \tag{3.11}$$

For the resulting within-class means, further denoted by $\hat{\eta}_i$, we have

$$\hat{\eta}_i = \mathbf{U}_{\mathbf{R}}^{'} \nu_i. \tag{3.12}$$

It has been proved in [21] that

$$(\mathbf{\Lambda_R})_{j,j} = 1, j = 1, \ldots, N_{\text{PCA}} - N_{\text{user}} + 1, \tag{3.13}$$

$$(\hat{\eta}_{\text{i}})_j = 0, j = 1, \ldots, N_{\text{PCA}} - N_{\text{user}} + 1. \tag{3.14}$$

This means that only the last $N_{\text{user}} - 1$ dimensions of $\mathbf{U}_{\mathbf{R}}^{'}\mathbf{R}$ can contribute to the verification. Therefore, a further dimension reduction is obtained by discarding the first $N_{\text{PCA}} - N_{\text{user}} + 1$ dimensions in $\mathbf{U}_{\mathbf{R}}^{'}\mathbf{R}$. This can be achieved by pre-multiplying $\mathbf{R}$ by a transformation matrix $\mathbf{U}_{\text{LDA}}$, with $\mathbf{U}_{\text{LDA}}$ the submatrix of $\mathbf{U}_{\mathbf{R}}$ consisting of the last $N_{\text{user}} - 1$ columns. The subscript LDA which stands for Linear Discriminant Analysis is used, as this operation is in fact a dimension reduction by means of LDA [39].

The sequence of transformations described above can be denoted as an $(N_{\text{user}} - 1) \times N_{\text{raw}}$ matrix

$$\mathbf{F} = \sqrt{N_{\text{ex}} - 1}\mathbf{U}_{\text{LDA}}^{'}\mathbf{S}_{\text{PCA}}^{-1}\mathbf{U}_{\text{PCA}}^{'}. \tag{3.15}$$

Let $\hat{\mathbf{z}} = \mathbf{F}\mathbf{z}$ denote the transformed feature vector of a measured image $\mathbf{z}$, then the matching score (3.4) computed for class c becomes

$$M(\hat{\mathbf{z}}) = -\frac{1}{2}(\hat{\mathbf{z}} - \hat{\mu}_{\text{c}})^{'}\mathbf{\Lambda}_{\mathbf{R}}^{-1}(\hat{\mathbf{z}} - \hat{\mu}_{\text{c}}) + \frac{1}{2}\hat{\mathbf{z}}^{'}\hat{\mathbf{z}} - \frac{1}{2}\log|\mathbf{\Lambda}_{\mathbf{R}}|. \tag{3.16}$$

Note that the derivation of (3.16) is based on the assumption that $\mathbf{X}$ has zero column mean. It can be proved that in the general case, whether or not the column mean of $\mathbf{X}$ is zero, the matching score for class c is

$$\begin{aligned} M(\hat{\mathbf{z}}) \;=\; & -\frac{1}{2}(\hat{\mathbf{z}} - \hat{\mu}_{\text{c}})^{'}\mathbf{\Lambda}_{\text{R}}^{-1}(\hat{\mathbf{z}} - \hat{\mu}_{\text{c}}) \\ & + \frac{1}{2}(\hat{\mathbf{z}} - \hat{\mu}_{\text{T}})^{'}(\hat{\mathbf{z}} - \hat{\mu}_{\text{T}}) \\ & - \frac{1}{2}\log|\mathbf{\Lambda}_{\mathbf{R}}|, \end{aligned} \tag{3.17}$$

where

$$\hat{\mathbf{z}} = \mathbf{F}\mathbf{z}, \tag{3.18}$$

$$\hat{\mu}_{\text{c}} = \mathbf{F}\mu_{\text{c}}, \tag{3.19}$$

$$\hat{\mu}_{\text{T}} = \mathbf{F}\mu_{\text{T}}. \tag{3.20}$$

That is, to compute the matching score a total of four entities have to be estimated from the training data: μ_{c}, the local mean vector of class

c before transformation; μ_{T}, the total mean vector before transformation; $\mathbf{F}$, the transformation matrix; and $\mathbf{\Lambda_R}$, the total covariance matrix after transformation.

3.3 Experiments, results and discussion

In this section we present the experimental results of grip-pattern verification, using the likelihood-ratio classifier as described in Section 3.2. These results will then be analyzed in terms of both the characteristics of the grip-pattern data, and the property of the likelihood-ratio classifier. Based on the analysis, solutions will be proposed to improve the verification performance.

Experiment set-up

We did two types of experiments: the within-session experiment, where the grip patterns for training and testing were clearly separated, but came from the same collection session; and, the across-session experiment, where the grip patterns collected in two different sessions were used for training and testing, respectively.

The verification performance was evaluated by the overall equal-error rate of all the subjects. The equal-error rate is the value of the false-acceptance rate, when the verification system is tuned in such a way that the false-acceptance rate and the false-rejection rate are equal. Note that our research goal is, however, to minimize the false-acceptance rate of the verification system at the false-rejection rate equal to 10^{-4}, as described in Section 1.2. The reason that we selected the equal-error rate as a measure of the verification performance, instead, was twofold. First, for a long period of our research we mainly focused on improving the verification performance of the system in general. Second, the equal-error rate is a commonly used measure of the performance of a biometric recognition system.

In the across-session experiment, the overall equal-error rate was estimated from the matching scores, as expressed in (3.17), of all the genuine users and impostors. In the within-session experiment, the overall equal-error rate was estimated based on all the matching scores, obtained from

20 runs. In each single run, 75% of the grip patterns were randomly chosen for training, and the remaining 25% for testing.

Prior to classification we processed the grip-pattern data in two steps, which proved to be beneficial to the verification results. First, each grip-pattern image was scaled such that the values of all pixels were in the range $[0, 1]$. This was to avoid that the features in greater numeric ranges dominated those in smaller numeric ranges [40]. Second, we applied a logarithm transformation on each scaled image. In this way, the contrast between a grip pattern and its background in an image was greatly enhanced, as shown in Figure 3.1.

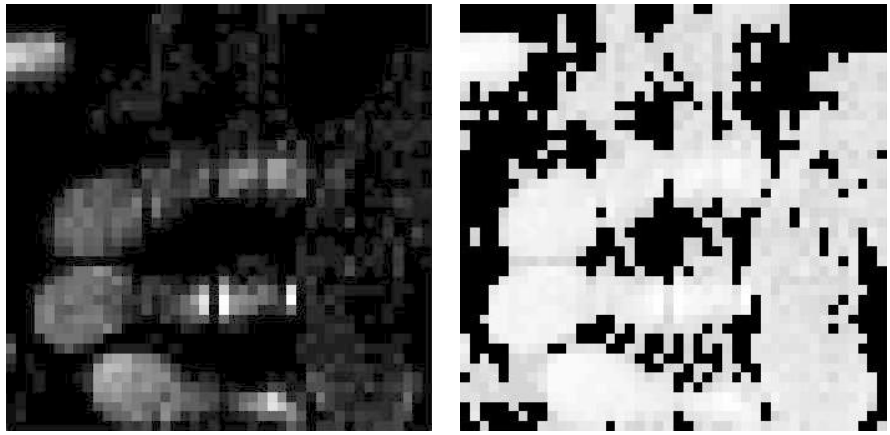

(a) Before (b) After

Figure 3.1: A grip-pattern image before and after logarithm transformation.

Experimental results

Before the verification performance can be assessed, both the feature dimensions kept after PCA and LDA have to be set. It was found that the verification performance was not sensitive to the feature dimension after PCA. As a (flat) optimum we selected $N_{\text{PCA}} = 3N_{\text{user}}$, with N_{user} the number of subjects in the training set. It was described in Section 3.2 that after LDA, only the last $N_{\text{user}} - 1$ dimensions of data can contribute to

the verification. Yet, as one can imagine, it is possible that the verification performance becomes even better if the feature dimension is further reduced. We found, however, that the verification performance only became worse when further dimension reduction was applied. Figure 3.2, for example, shows the equal-error rate as a function of feature dimension after LDA. Here, the grip patterns for training and testing are those recorded in the first and second collection sessions, respectively. As a result, in our experiment the feature dimension after LDA was reduced to $N_{\text{user}} - 1$.

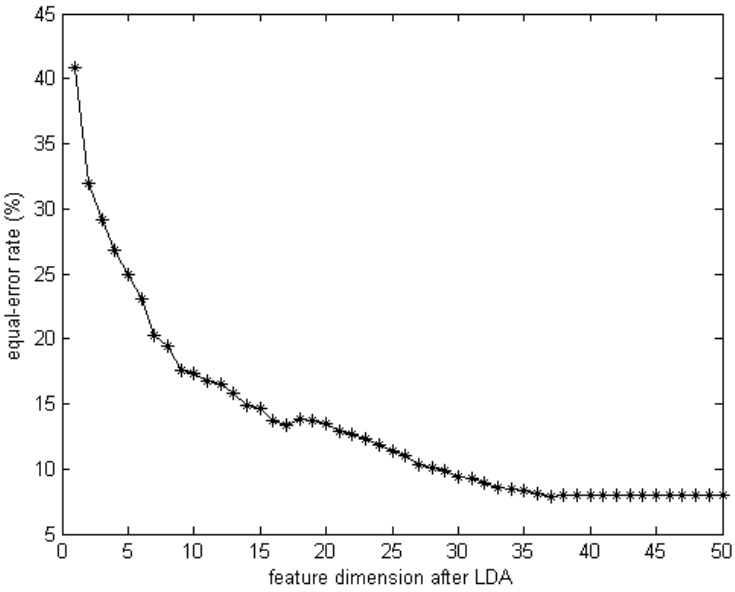

Figure 3.2: Equal-error rate as a function of feature dimension after LDA.

The experimental results for within-session verification are presented in Table 3.1. As a reference the verification result based on grip patterns recorded from a group of untrained subjects previously, as described in Section 1.2, is also presented here, denoted as Session 0. Table 3.2 shows the experimental results for across-session verification. One can see from Figure 3.3 the false-acceptance and false-rejection rate curves, of both within-

Table 3.1: Within-session verification results.

Session	Equal-error rate (%)
0	1.4
1	0.5
2	0.8
3	0.4

Table 3.2: Across-session verification results.

Train session	Test session	Equal-error rate (%)
2	1	5.5
3	1	14.7
1	2	7.9
3	2	20.2
1	3	24.1
2	3	19.0

session and across-session verification. As an example, the grip patterns recorded in the first and second sessions are used for training and testing, respectively.

The experimental results indicate that when the grip-pattern data for training and testing came from the same session, the verification performance was fairly good. Also, one can see that the verification performance was much better with the grip patterns recorded from the police officers, than from the untrained subjects. However, the verification performance became much worse when the grip patterns for training and testing were recorded in two different sessions, respectively. Since in practice there is always a time lapse between the data enrollment and verification, the across-session verification performance is more relevant and, therefore, has to be improved.

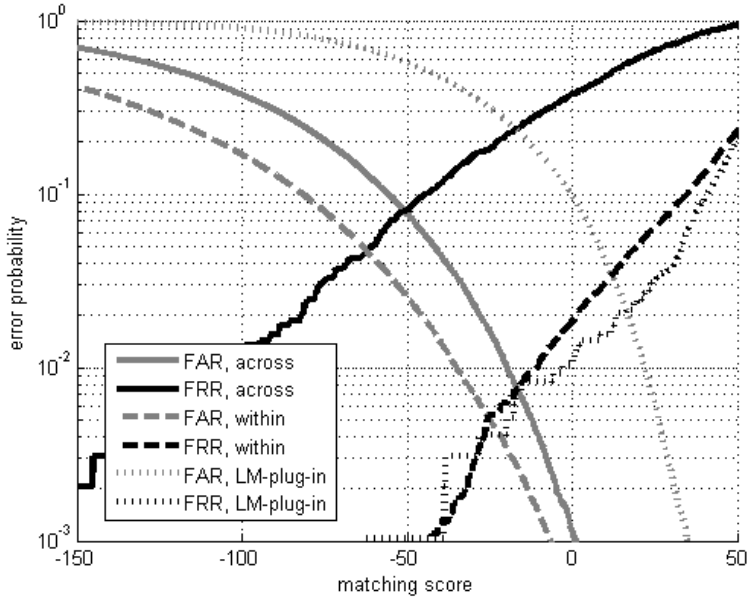

Figure 3.3: False-acceptance and false-rejection rate curves of within-session, across-session and LM-plug-in experiments.

Discussion

Data characteristics

Comparing the grip-pattern images of the same subjects recorded in different sessions, we observed large across-session variations. That is, data drift occurred for grip patterns. Note that this was different from our expectation that the grip patterns of trained users are very stable. Specifically, two types of variations were observed. First, the pressure distribution of a subject's grip-pattern image recorded in one session was usually very different, from his or her grip-pattern image recorded in another session. Second, for some subjects the horizontal or vertical shift of hand position was found from one session to another. However, we also found that the hand shape of a subject remained constant across sessions. This is easy to understand, as the hand shape of a subject is a physical characteristic, which does not

change so rapidly. The stability of one's hand shape is also supported by the fact that hand geometry has proved to perform reasonably well for identity verification [12], [13], [14], [15] and [16].

The characteristics of grip-pattern images described above can be illustrated by Figure 3.4, where the two images are from the same subject yet recorded in different sessions. On the one hand, one can see that these two images have very different pressure distributions. And, the grip pattern in Figure 3.4(b) is located higher than that in Figure 3.4(a). On the other hand, the hand shape of the subject does not change much from one image to another.

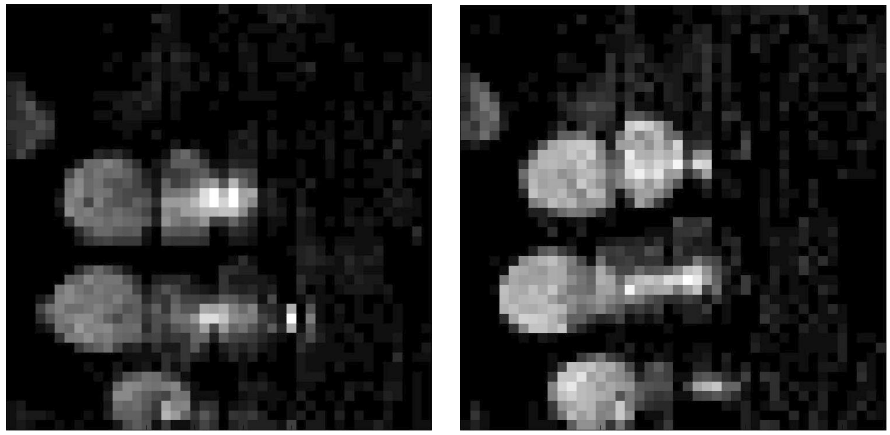

(a) (b)

Figure 3.4: Grip-pattern images of a subject in different sessions.

As shown in (3.17), four entities have to be estimated from the training data to compute the matching score between a measurement and a class c: μ_c, the local mean vector of class c before transformation; μ_T, the total mean vector before transformation; $\mathbf{F}$, the transformation matrix; and $\mathbf{\Lambda_R}$, the total covariance matrix after transformation. As a result of the across-session variations of grip patterns, the value of each entity varied from one session to another. In order to find out the variation of which entity degraded the verification performance the most, we did the following "plug-in" experiment. First, we randomly split the test set into two subsets of

Table 3.3: Across-session verification results in equal-error rate (%) with one entity estimated from subset $\mathbf{D}_2$.

Train session	2	3	1	3	1	2
Test session	1	1	2	2	3	3
RF	5.5	14.7	7.9	20.2	24.1	19.0
μ_c	**1.0**	**2.2**	**2.1**	**2.8**	**2.7**	**2.6**
$\mathbf{\Lambda_R}$	6.0	13.5	7.3	17.3	24.3	19.9
μ_T	5.5	14.9	8.0	20.0	24.0	19.2
$\mathbf{F}$	3.8	23.8	3.2	18.4	13.7	14.9

equal size, namely, $\mathbf{D}_1$ and $\mathbf{D}_2$; then we used one subset, for example, $\mathbf{D}_1$ for testing. In computation of (3.17), each time we estimated three out of the four entities from the training set, yet the fourth one from subset $\mathbf{D}_2$.

The verification results for the "plug-in" experiment are presented in Table 3.3. The last four rows show the equal-error rates, with μ_c, $\mathbf{\Lambda_R}$, μ_T, and $\mathbf{F}$ estimated from subset $\mathbf{D}_2$, respectively. As a reference, the experimental results in Table 3.2 are also presented, denoted as 'RF'. One can see that the verification performance was improved dramatically, when μ_c was estimated from subset $\mathbf{D}_2$. In this case the verification results even became close to those in the within-session experiment, as shown in Table 3.1. This indicates, therefore, that the verification performance was degraded the most by the variation of the mean value of a subject's grip patterns across sessions.

The false-acceptance and false-rejection rate curves, after "plugging in" the mean values, are shown in Figure 3.3. As earlier, the grip patterns recorded in the first and second sessions are used for training and testing, respectively. One can see that the false-rejection rate has decreased compared to the case of across-session verification. Yet, the false-acceptance rate has increased. However, this effect on the equal-error rate is not as strong as that of the decrease of the false-rejection rate.

Algorithm property

Besides the across-session variations of the grip patterns, the unsatisfactory verification performance was also due to a property of the likelihood-ratio

classifier. As described in Section 3.2, this classifier requires estimating the probability density function of the grip-pattern data, from a set of training samples. The verification performance, therefore, depends largely on how similar the estimate from the training data is to the actual situation of the test data. That is, this type of classifier does not have a good generalization property, and is not robust enough to the problem of data drift. If large variations occur between the data for training and testing, the verification performance will be degraded significantly.

Possible strategies to improve the performance

The verification performance may be improved in three ways, given the characteristics of the grip-pattern data and the property of the likelihood-ratio classifier. First, we may model the across-session variations of the grip-patterns images, during training of the classifier. The method to achieve this is called double-trained model, and will be described in Section 3.3. Or, we may reduce the variations by applying some image processing techniques prior to classification. For example, we may apply techniques that equalize the local pressure values in an image, to reduce the difference in pressure distribution between two images of the same subject. Also, image registration methods may help align two grip-pattern images with hand shifts. These techniques will be described in Chapter 4 and 5 of this thesis.

Second, instead of using the low-level features, i.e. the pressure values measured as the output of the sensor, we may also investigate the verification performance using high-level features extracted from grip-pattern images, i.e. the physical characteristics of a hand or hand pressure. This is mainly inspired by our observation that the hand shape remains constant for the same subject across sessions. Grip-pattern verification based on high-level features will be done and analyzed in Chapter 8.

Third, we may turn to some other type of classifier. The Support Vector Machine classifier seems a promising choice, attributed to its good generalization property [32], [33]. As a contrast to a probability-density-based classifier, the support vector machine classifier does not estimate the probability density function of the data. Instead, it maximizes the margin between different classes and has been proved to be more robust to the data drift in many cases [34], [35], [36], [37], [38]. The performance of grip-pattern

Table 3.4: Across-session verification results with DTM applied.

Train session	Test session	Equal-error rate (%)
2+3	1	4.0
1+3	2	5.7
1+2	3	13.7

verification using support vector machine classifier will be investigated in Chapter 9.

Double-trained model

According to the characteristics of grip-pattern data, the verification performance may be improved by modelling their across-session variations during training of the classifier. Therefore, we applied the double-trained model (DTM). Specifically, we combined the grip patterns recorded in two out of three collection sessions for training, and used those of the remaining session for testing. In this way, both the variation of pressure distribution and that of hand position were modelled much better in the training procedure, compared to the case where the classifier was trained based on grip patterns collected in one session alone.

Table 3.4 presents the experimental results. Comparing them to those shown in Table 3.2, one can see that the verification performance has been improved greatly with the application of DTM.

3.4 Conclusions

Grip-pattern verification was done using a likelihood-ratio classifier. It has been shown that the grip patterns contain useful information for identity verification. However, the verification performance was not good. This was mainly due to the variations of pressure distribution and hand position between the training and test images of a subject. Further analysis shows that it was the variation of the mean value of a subject's grip patterns, that degraded the verification performance the most. And, since the likelihood-ratio classifier is probability density based, it is not robust to the problem

of data drift. Nonetheless, we also found that the hand shape of a subject remained constant, in both the training and test images.

Given the analysis, three solutions were proposed to improve the verification performance. First, we may model the data variations during training of the classifier, or reduce the data variations with some image processing techniques. Second, verification may be done with high-level features, i.e. the physical characteristics of a hand or hand pressure. Third, we may use some other type of classifier, which is not based on the data probability density.

In this chapter, we applied the method of double-trained model (DTM), where the grip patterns collected from two sessions were combined for training. With DTM the verification performance was greatly improved, since the across-session variations were modelled much better, than the case where the classifier was trained based on the data recorded in one collection session alone. The other approaches to improve the verification performance will be described later in this thesis.

Chapter 4

Registration of grip-pattern images[1]

Abstract. *In this chapter two registration methods, template-matching registration and maximum-matching-score registration, are proposed to reduce the variations of hand positions between the probe image and the gallery image of a subject. The experimental results based on these approaches are compared. Further, a fused classifier is applied, using discriminative information of the grip-pattern data obtained with the application of both registration methods, which significantly improves the results.*

4.1 Introduction

As described in Chapter 3, after analyzing the images collected in different sessions, we have found that even though the grip-pattern images from a certain subject collected in one session look fairly similar, a subject tends to produce grip-pattern data with larger variations across sessions. First, a variation of pressure distributions occurs between grip patterns of a subject across sessions. Second, another type of variation results from the hand shift of a subject. These variations are illustrated in Figure 4.1. Therefore,

[1]This chapter is based on [26].

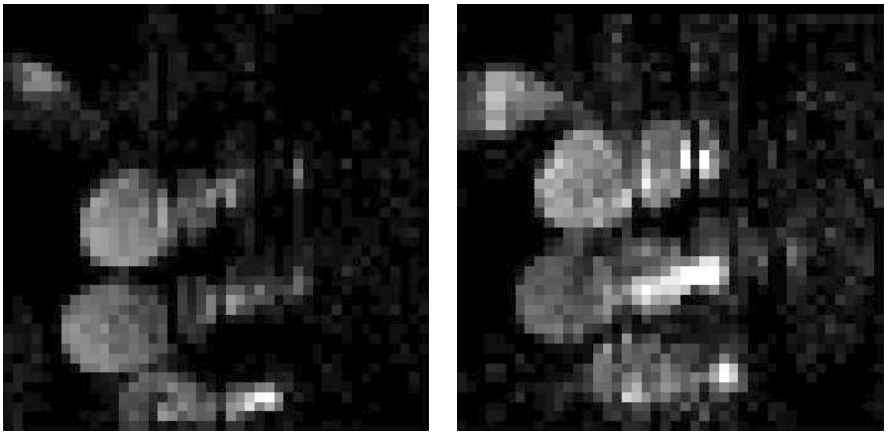

(a) (b)

Figure 4.1: Grip-pattern images of a subject in different sessions.

the verification results may be improved by reducing the data variations across sessions.

In order to reduce the variation caused by the hand shift, we have applied image registration methods for aligning a test image to a registration template. Rotation or scaling of a measured image with respect to a registration template has not been observed, and is not very likely to occur regarding the device at hand. Therefore, in our experiments we only used shifts for aligning the grip patterns. Specifically, two types of registration methods were implemented as a preprocessing step prior to classification respectively, namely, template-matching registration (TMR) [41] and maximum-matching-score registration (MMSR) [42], [43]. The reason that these two techniques were investigated is that TMR is a standard method for registration, while MMSR is a promising new method.

It has been found that TMR is able to effectively improve the across-session verification performance, whereas MMSR is not. However, the hand shift measured by MMSR has proved particularly useful in distinguishing impostors from genuine users. If two images belong to the same subject, the hand shift value produced by MMSR is on average much smaller than if they belong to different subjects. Inspired by both this observation and the

concept of fusion of classifiers [44], [45], we designed a new, fused, classifier based on both the grip pattern and the hand shift. This has further reduced the verification error rates significantly.

This chapter presents and compares the verification results with data preprocessed by TMR and MMSR prior to classification, respectively. The remainder of this chapter is organized as follows. Section 4.2 briefly describes the registration approaches. Next, the experimental results using the different registration methods will be compared in Section 4.3. Subsequently, Section 4.4 presents the verification results by using discriminative information of the grip-pattern data, obtained with the application of both registration methods. Finally, conclusions will be given in Section 4.5.

4.2 Registration method description

In TMR, the normalized cross correlation of a measured image and a registration-template image is computed. The location of the pixel with the highest value in the output image determines the hand shift value of the measured image with respect to the template image. If the measured image is well aligned to the template image, this pixel should be located precisely at the origin of the output image. In our case, the measured image, the template image, and the output image are all of the same size. The shifted version of an original image after TMR can be described as

$$\mathbf{z}_0 = \arg\max_{\tilde{\mathbf{z}}} \frac{\mathbf{y}' \cdot \tilde{\mathbf{z}}}{\|\mathbf{y}\|\|\tilde{\mathbf{z}}\|}, \tag{4.1}$$

where $\tilde{\mathbf{z}}$ denotes a shifted version of an original image $\mathbf{z}$, and $\mathbf{y}$ denotes the registration-template image. The symbol $'$ denotes vector or matrix transposition. Among all the training samples of a certain subject, the one with the minimal Euclidean distance to the mean value of this subject was used as the registration template of this subject.

In MMSR, a measured image is aligned such that the matching score, $M(\mathbf{z})$ in (3.4), attains its maximum. Specifically, an image is shifted pixel by pixel in both the horizontal and vertical directions. After each shift, a new matching score is computed. This procedure continues until the original image has been shifted to all the possible locations within a predefined

scope of 20 pixels in each direction. In the end, the shifted image with the maximum matching score is selected as the registration result. It can be represented as

$$\mathbf{z}_0 = \arg\max_{\tilde{\mathbf{z}}} \mathrm{M}(\tilde{\mathbf{z}}), \tag{4.2}$$

where $\tilde{\mathbf{z}}$ denotes a shifted version of the original image $\mathbf{z}$. Note that in TMR the "template image" refers to the registration template image, whereas in MMSR it refers to the recognition template image.

4.3 Experimental results

Across-session experiments were done using the likelihood-ratio classifier. The verification performance was evaluated by the overall equal-error rate of all the subjects. It was estimated from the likelihood ratios of all the genuine users and impostors, as described in Section 3.3.

During training the data registration was done in two steps. First, we did user-nonspecific registration by TMR to align all the training samples to their mean image. Second, TMR was applied to the training data to build up a stable after-registration model with user-specific registration templates. Specifically, among all the training samples of a certain subject, the one with the minimal Euclidean distance to the mean value of this subject was used as the registration template of this subject. All the other training samples were then aligned to this image. This procedure was repeated iteratively until no more shift occurred for each image. During testing prior to classification we applied user-specific registration to the test data by TMR and MMSR, respectively. We also did the experiment where MMSR, in stead of TMR, was applied to the training data in order to obtain a registered training set. The verification performance was, however, much worse than that in the case where TMR was used.

We only used the lower-left part, of size 33×33, of each image, where the fingers of the subjects are located, for computing the cross correlation. See Figure 4.2. There are two reasons for this. First, sometimes the positions of the thumb and fingers do not always change in the same way (see Figure 4.1). Second, according to our observation, sometimes the pressure

Table 4.1: Across-session verification results in equal-error rate (%) with and without registration approaches.

Train session	2	3	1	3	1	2
Test session	1	1	2	2	3	3
RF	5.5	14.7	7.9	20.2	24.1	19.0
TMR	**3.9**	**12.9**	**6.0**	**17.8**	**18.4**	**18.9**
MMSR	5.8	17.7	8.0	22.9	27.7	22.8

pattern of one's thumb is rather unclear or not even present, and therefore, not reliable enough for the registration.

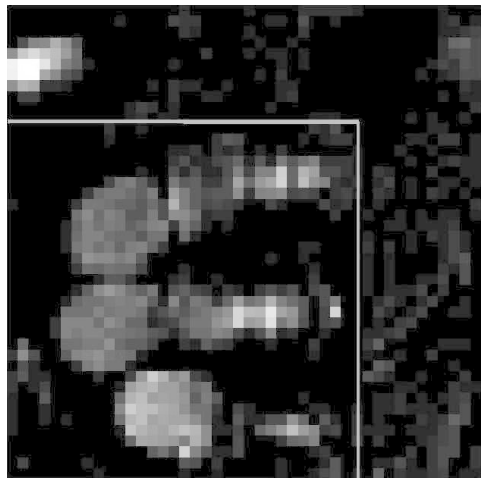

Figure 4.2: Only the lower-left part of an image, where fingers of a subject are located, is used for template-matching registration.

Tables 4.1 presents the experimental results. As a reference, the results without any data registration are also shown, represented as RF. One can see that the results have been improved when TMR is applied to the test data, whereas the results have become worse when the test data are preprocessed by MMSR.

The corresponding false-acceptance and false-rejection rate curves can be found in Figure 4.3, where the grip-pattern data from the second and

first sessions were used for training and testing, respectively. One can see that when either of the two registration methods is in use, for a certain threshold of the matching score, the false-rejection rate decreases and the false-acceptance rate increases compared to their counterparts without any registration step. However, in the case of TMR, the false-rejection rate decreases more than the false-acceptance rate increases, whereas in the case of MMSR it is the other way around.

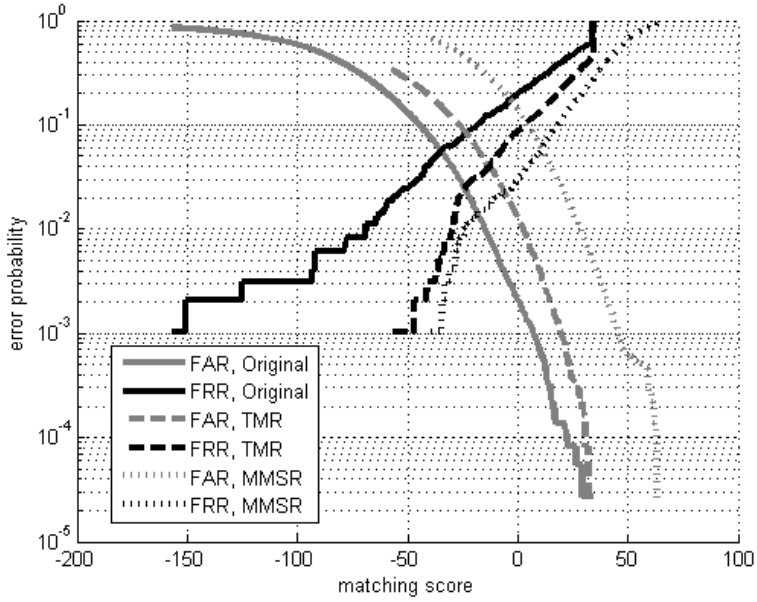

Figure 4.3: Comparison of false-acceptance and false-rejection rate curves obtained in different conditions.

The different effects of these registration methods result from their working principles and the characteristics of the grip-pattern images. Apart from the hand shift, a large variation of the pressure distribution may exist between a measured grip-pattern image and the template to which it is compared (see Figure 4.1). Therefore, neither of the registration methods may yield an ideal result. The increase in equal-error rate when MMSR is applied, shown in Table 4.1, may be explained as follows. Since the original

matching scores of the impostors will be relatively low compared to those of the genuine users, the increase in the matching scores of the impostors will be on average larger than of the genuine users. That is, the effect of the increasing false-acceptance rate will be stronger than that of the decreasing false-rejection rate. In contrast, TMR, not aiming at a maximum matching score, does not increase the false-acceptance rate as much as MMSR does. As a net effect, TMR improves the verification results, whereas MMSR does not.

4.4 Classifier based on both grip pattern and hand shift

For each measured image, the application of both TMR and MMSR results in a value of hand shift. We found that if the measured image and the template image belong to the same subject, the produced hand shift value is on average much smaller than if they belong to two different subjects, respectively. This is easy to understand, since the variations of grip-pattern data from two different subjects are supposed to be larger than those from the same subject. Therefore, it is very likely that the grip pattern of the impostor is shifted more than that of the genuine user. Also, we found that the genuine and impostor hand shifts are more discriminative if they are produced by MMSR rather than by TMR. Specifically, the hand shifts from the impostors produced by MMSR are on average larger than those produced by TMR, whereas the hand shifts from the genuine users have similar values produced by these two approaches.

The characteristics of hand shifts produced by TMR and MMSR, as mentioned above, can be illustrated by Figure 4.4 and Figure 4.5. They present the probability distributions of the l^2-norm of the hand shifts in both the vertical and horizontal directions as measured by TMR and MMSR, respectively. The training data are from the third session, and the test data are from the first session. One can see that the hand shifts from the impostors are mostly of small values in Figure 4.4, yet those in Figure 4.5 are of much greater values in general. Table 4.2 lists the means and standard deviations of l^2-norm of the hand shifts in both the vertical and horizontal directions as measured by TMR and MMSR, respectively. The 'G' denotes

Table 4.2: Means and Standard deviations of l^2-norm of hand shifts in both vertical and horizontal directions as measured by TMR and MMSR.

	Mean	Standard deviation
G_{TMR}	1.6	2.0
G_{MMSR}	1.6	1.4
I_{TMR}	3.4	3.0
I_{MMSR}	5.9	2.9

'genuine user' and 'I' denotes 'imposter'.

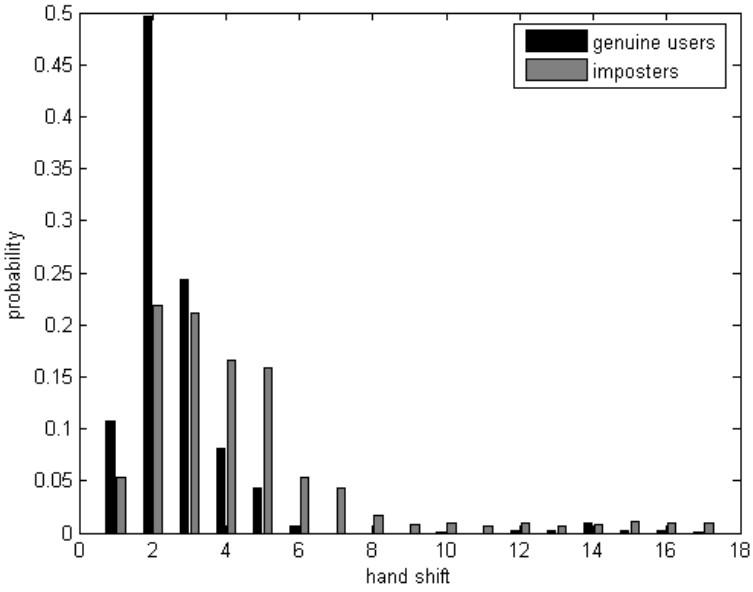

Figure 4.4: Probability distributions of hand shift after TMR.

The reason that the hand shifts from the impostors obtained by TMR are much smaller in general than those obtained by MMSR, may be because the registration results of TMR depend more on the global shapes of the grip-pattern images, which constrains the shift values of the images from

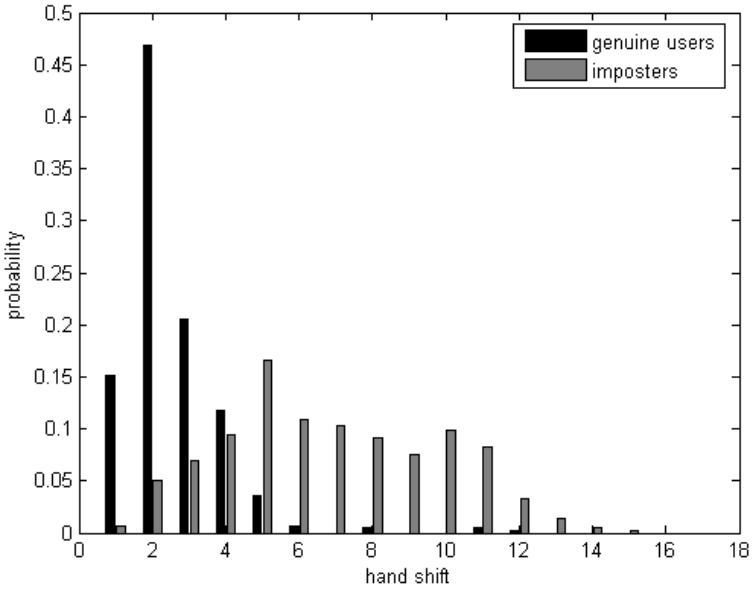

Figure 4.5: Probability distributions of hand shift after MMSR.

impostors to a relatively small value. See Figure 4.6 for an illustration of this effect.

Inspired by the special characteristic of the hand shift produced by MMSR, we implemented a new classifier as a combination of two other classifiers. Specifically, one is based on grip patterns using the likelihood-ratio classifier, with TMR as a preprocessing step. The other one performs verification based on the minus l^2-norm of hand shifts produced by MMSR. Note that in both classifiers, TMR is applied to the training data to build up a stable after-registration model. A measured grip-pattern image is verified as being from the genuine user if, and only if, the verification results given by both classifiers are positive. This is similar to the application of the threshold-optimized AND-rule fusion, described in [46], [47], [48] and [49]. The difference is that we did not optimize the threshold values as in the case of threshold-optimized AND-rule fusion. Instead, regarding a

Table 4.3: Across-session verification results in equal-error rate (%) with reference and improved classifiers.

Train session	2	3	1	3	1	2
Test session	1	1	2	2	3	3
RF	5.5	14.7	7.9	20.2	24.1	19.0
TMR	3.9	12.9	6.0	17.8	18.4	18.9
COM	**3.4**	**8.1**	**5.0**	**12.4**	**11.0**	**12.6**

threshold value of one classifier, the threshold of the other classifier was set as *any* of the possible values, for making a verification decision.

The process of classification can be illustrated by Figure 4.7, where circles and stars represent images from the genuine users and the impostors, respectively. Let the threshold value be set as -5 on the horizontal axis, and as -100 on the vertical axis. Then only those corresponding images, represented as scatter points located in the upper-right corner on the plot, will be classified as being from the genuine user, whereas all the other images will be rejected.

Table 4.3 shows that the verification results have been further improved by using the combination of these two classifiers, denoted as COM. For easy comparison, the experimental results in the third and the fourth rows of Table 4.1 are presented as well.

4.5 Conclusions

In order to improve the verification results, two registration approaches are applied to reduce the across-session variation caused by the hand shift: template-matching registration (TMR) and maximum-matching score registration (MMSR). Both methods prove to be useful to improve the verification performance of the system. Specifically, verification error rates are reduced effectively when the TMR is used as a preprocessing step prior to classification. Also, the hand shift value produced by the MMSR is on average much smaller if the measured image and the template image belong to the same subject rather than if they belong to two different subjects, respectively. As a result, a new classifier is designed based on both

the grip patterns preprocessed by TMR, and the hand shifts produced by MMSR. The application of this additional classifier has greatly improved the verification results further.

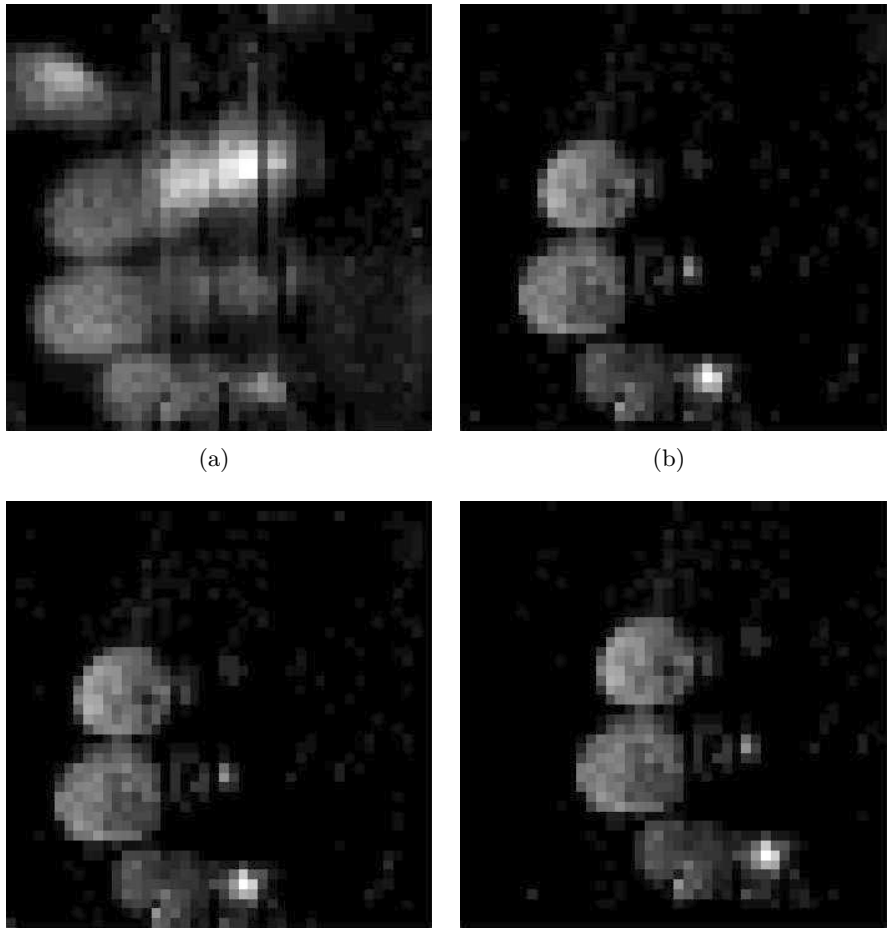

(a) (b)

(c) (d)

Figure 4.6: (a) Template image of one subject from the first session. (b) Image of an imposter from the second session. (c) Image of the imposter after TMR. (d) Image of the imposter after MMSR.

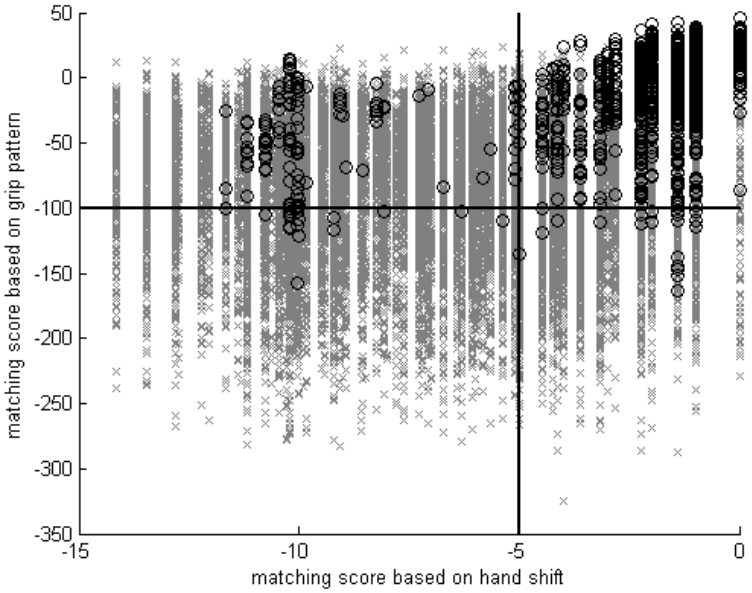

Figure 4.7: Scatter graph of matching scores computed from both grip patterns and hand shifts.

Chapter 5

Local absolute binary patterns as image preprocessing for grip-pattern recognition[1]

Abstract. *In this chapter we propose a novel preprocessing technique for grip-pattern classification, namely, Local Absolute Binary Patterns. With respect to a certain pixel in an image, Local Absolute Binary Patterns processing quantifies how its neighboring pixels fluctuate. It will be shown that this technique can reduce the variation of pressure distribution between the probe and gallery images of a subject. Also, information of the hand shape in an image can be extracted. As a result, a significant improvement of the verification performance will be achieved.*

5.1 Introduction

As described in Chapter 3, after analyzing the images collected in different sessions we found that a subject tends to produce the grip patterns with large variations across sessions, even though his or her grip patterns collected within one session are fairly similar. Basically, two types of variations

[1]This chapter is based on [25].

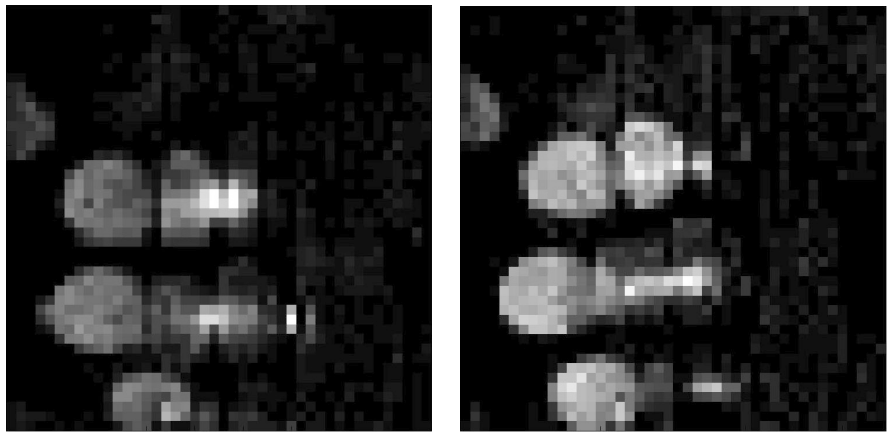

(a) (b)

Figure 5.1: Grip-pattern images of a subject in different sessions.

were observed. First, a variation of pressure distributions occurs between the grip patterns of a subject across sessions. Second, another type of variation results from the hand shift of a subject. However, despite the variations of grip patterns, we also found that the hand shape, as it occurs in one's grip-pattern image, is almost constant across sessions. This is easy to understand, as the hand shape of a subject is a physical characteristic, which does not change so rapidly.

Figure 5.1 shows two images collected from one subject in different sessions. One can see that the pressure distributions in these two images are quite different. Also, the grip pattern in Figure 5.1(b) is located higher, on average, than that in Figure 5.1(a). However, the hand shapes of the grip patterns in these two images look quite similar.

Based on characteristics of the grip-pattern images, the verification performance may be improved by reducing the across-session variations of grip patterns, or by extracting information of hand shapes in the images. In earlier work, we used template-matching registration to reduce the variation caused by the hand shift. Also, we applied a double-trained model, with which the data from two out of three sessions were combined for training, and those of the remaining session were used for testing. By doing this,

the data variations across sessions were much better modelled during training procedure of the classifier, compared to the case where the grip-pattern data from only one collection session were used for training. As a result, the verification performance has been greatly improved with both the application of the template-matching registration, and that of the double-trained model.

In this chapter we propose a preprocessing approach for grip-patter classification, namely, Local Absolute Binary Patterns (LABP). This technique is not only capable of reducing the across-session variation of the pressure distribution in the grip-pattern images, but also, it can extract information of hand shapes in the images by enhancing the contrast between the grip patterns and the background. Specifically, with respect to a certain pixel in an image, the LABP processing quantifies how its neighboring pixels fluctuate. It was found that the application of LABP can further improve the verification performance significantly.

For comparison, the experimental results with the grip patterns preprocessed by a type of high-pass filter will also be presented. The reason for this is twofold. First, as discussed earlier, the performance of grip-pattern verification may be improved by extracting information of hand shapes in the images. Second, it can be proved that the hand part and the background in a grip-pattern image can be regarded as the high-frequency and low-frequency component of this image, respectively. That is, the application of a high-pass filter may improve the verification performance by enhancing the contrast between a grip pattern and its background, as in the case of LABP processing.

This chapter presents and analyzes the verification results, with the grip-pattern images preprocessed by LABP processing prior to classification. The remainder of this chapter is organized as follows: Section 5.2 describes the approach of LABP processing. Next, Section 5.3 compares it to a type of high-pass filter, generated from the Gaussian low-pass filter. The experimental results will be presented and discussed in Section 5.4. Finally, conclusions will be given in Section 5.5.

5.2 Local Absolute Binary Patterns

The development of Local Absolute Binary Patterns (LABP) is mainly based on two observations of the grip-pattern images. First, regarding a certain pixel located within the hand part, the values of its neighboring pixels fluctuate to a large degree on average. This may be because the grip pattern, by nature, is a type of locally fluctuating data. Second, in contrast, the pressure distribution in areas within the background is much smoother, as most pixels in the background have a pressure value zero. Figure 5.2, as an example, illustrates these observations.

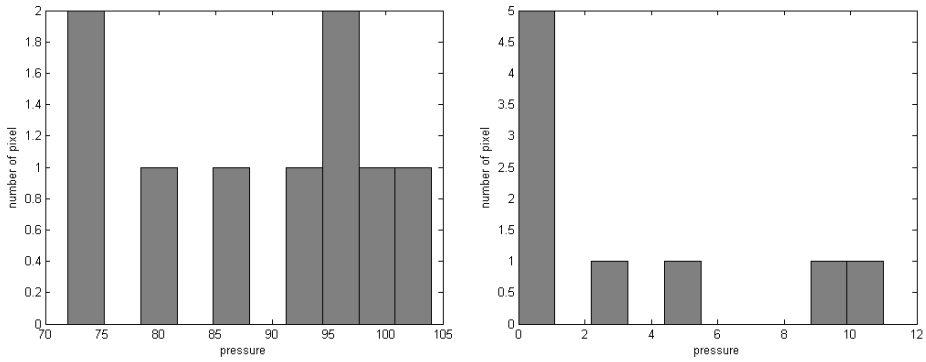

(a) (b)

Figure 5.2: (a) Histogram of pixel values in a three-by-three neighborhood within the hand part of the grip-pattern image in Figure 5.1(a). (b) Histogram of pixel values in a three-by-three neighborhood within the background of the grip-pattern image in Figure 5.1(a).

Inspired by these special characteristics of grip-pattern images, and based on previous work on local binary patterns (LBP) [50], [51], we have developed LABP. With respect to each certain pixel in an image, the LABP processing quantifies how its eight neighboring pixels fluctuate. It is performed in two steps. First, for a given center pixel (x_c, y_c), the values of the neighboring pixels are compared to that of the center pixel (x_c, y_c), one by one. As a result of each comparison, either a 1 or a 0 is recorded: if the absolute difference between two pixels is greater than a threshold,

31	53	42
14	39	46
26	33	42

1	1	0
1		1
1	1	0

Figure 5.3: Top: a three-by-three neighborhood in an image before LABP. Bottom: comparison results between center pixel and neighboring pixels.

the comparison result will be 1; otherwise, the comparison result will be 0. In this way, eight comparison results are computed. Second, the value of the center pixel is replaced by the sum of these eight comparison results. Therefore, the new value of (x_c, y_c) will be an integer between 0 and 8. In this way each pixel in an image is assigned a new value.

The resulting value of the center pixel can be expressed as follows:

$$\mathrm{LABP}(x_c, y_c, t) = \sum_{m=0}^{7} \mathrm{u}(|i_c - i_m| - t), \tag{5.1}$$

where i_c and i_m represent the values of the center pixel (x_c, y_c) and one of its neighboring pixels, respectively. The threshold t is a positive constant and needs to be optimized. $\mathrm{u}(r)$ is the unit step function defined as:

$$\mathrm{u}(r) = \begin{cases} 1 & \text{if} \quad r \geq 0, \\ 0 & \text{if} \quad r < 0. \end{cases} \tag{5.2}$$

As an example, Figure 5.3 illustrates the procedure of LABP processing in a three-by-three neighborhood, supposing that the threshold value is $t = 5$. The resulting value of the center pixel after LABP is therefore 6.

In our experiment, both prior to and after LABP processing, each image was scaled as such that the values of all pixels were in the range $[0, 1]$. This was to avoid that the features in greater numeric ranges dominated those in smaller numeric ranges [40]. Also, after scaling the threshold value would be invariant to grip-pattern images with different global ranges of pixel values.

The grip-pattern images processed by LABP are shown in Figure 5.4(a) and Figure 5.4(b), for example, with t equal to 0.006. The original images are from the same subject collected in different sessions, as shown in Figure 5.1.

It has been shown in our experiment, that the performance of grip-pattern verification varied greatly with different values of threshold t applied. On the one hand, when the value of t was too small, the verification performance was degraded due to the noisy information in the image after LABP processing. On the other hand, when the value of t was too large, the local fluctuations of small values in the image were ignored during LABP processing. Figure 5.5 illustrates these two effects with t equal to 0 and 0.02 applied to Figure 5.1(a), respectively.

One can see from Figure 5.4 that LABP has mainly two effects on the original grip-pattern images. First, the across-session variation of pressure distribution has been greatly reduced. Further analysis shows that this is because after LABP processing, in each image the pressure values in different subareas within the hand part have become much more equalized, in comparison with the case before LABP processing. That is, in each image the local pressure differences have been significantly reduced. Therefore, by normalization of each image after LABP processing, the pressure difference between the corresponding subareas of two images from the same subject has been greatly reduced. The effect of pressure equalization between different subareas within the hand part of an image, is further illustrated in Figure 5.6. One can see that after LABP processing the pressure difference between neighborhood 1 and 2 is much smaller in general, than that before LABP processing.

Second, with LABP the information of the hand shape in an image has been extracted by enhancing the contrast between the grip pattern and its background. Comparing images before and after LABP processing, one can see that the grip pattern processed by LABP is much more distinctive with respect to the background. Note that after LABP processing the noise in some part of the background has been amplified too. The value of noise depends on that of the threshold t in (5.1). The importance of selecting a proper value for the threshold will be discussed in more detail in Section 5.4.

5.3 Alternative method: high-pass filter

Figure 5.2 shows that the hand part and the background in a grip-pattern image can be regarded as the high-frequency and low-frequency component of this image, respectively. Since the LABP processing enhances the contrast between the hand part and the background, in this sense it is comparable to a high-pass filter. In this section we briefly introduce a type of high-pass filter and compare it to the LABP processing. It is one of the mostly widely used high-pass filter, generated by applying a Gaussian low-pass filter to an image and then subtract the obtained low-frequency image from the original image [52].

A Gaussian filter is a weight kernel that approximates the profile of a Gaussian function along any row, column, or diagonal through the center. It is characterized by a standard deviation, expressed in terms of pixel dimensions, and calculated as

$$G(x, y, \sigma) = \frac{1}{2\pi\sigma^2} e^{-\frac{x^2+y^2}{2\sigma^2}} \tag{5.3}$$

where x and y are the distance in pixels from the center of the kernel. The standard deviation σ is the radius (in pixels) of the kernel, containing 68% of the integrated magnitude of the coefficients [52].

Figure 5.4(c) and Figure 5.4(d) show the grip-pattern images processed by the high-pass filter, with the original images being from the same subject collected in different sessions, as shown in Figure 5.1. Both prior to and after the high-pass filtering, as in the case of LABP processing, each image has been scaled as such that the values of all pixels are in the range $[0, 1]$. One can see that after the high-pass filtering, the grip pattern in each image has not become more distinctive in general with respect to its background as in the case of LABP processing, even though the *local* contrast within the hand part has been enhanced.

Also, as one can see from Figure 5.6, after the high-pass filtering the pressure differences between different subareas within the hand part of an image have not been reduced, in contrast to the case of LABP processing. Here the histograms of two three-by-three neighborhoods are shown, of Figure 5.1(a), Figure 5.4(a) and Figure 5.4(c), respectively. Note that Figure 5.4(a) shows the grip-pattern image in Figure 5.1(a) after LABP

processing, and Figure 5.4(c) shows that after the high-pass filtering.

5.4 Experiments, Results and Discussion

We did the grip-pattern verification based on a likelihood-ratio classifier. As described in Section 3.3, the verification performance was evaluated by the overall equal-error rate of all the subjects. It was estimated from the log-likelihood ratios of all the genuine users and impostors. In the experiments each image was preprocessed by either LABP, or the high-pass filter described in Section 5.3, prior to classification. Each image was scaled such that the values of all pixels were in the range $[0, 1]$, both prior to and after LABP processing or high-pass filtering. Also, we applied both template-matching registration and double-trained model in the experiments (see Section 5.1).

The procedure of searching for t_{opt}, the optimal value of the threshold t for LABP processing in (5.1), proved to be crucial as discussed in Section 5.2. Figure 5.7 shows the verification results with different values of t. Here, as an example, the grip patterns of the first and second collection sessions were used for training and testing, respectively. The searching scope of t_{opt} was chosen as $[0, 0.02]$. One can see that the verification results varied greatly with different values of t. Experiments were also done with different values of the standard deviation, i.e. σ in (5.3), selected for high-pass filtering, but it turned out that the verification results did not vary much.

We tried two methods on data in the training set to find out t_{opt} for LABP processing. In the first method, with each value of t, the overall equal-error rate was estimated based on all the matching scores obtained from 20 runs. In each single run, 75% of the data in the training set were randomly chosen for training, and the remaining 25% for testing. The value of t corresponding to the lowest equal-error rate was then selected as t_{opt}. Finally, with $t = t_{\text{opt}}$, the data of the whole training set were used to train the classifier.

In the second method, data in the training set were divided into two parts, according to in which session they were collected. Since the double-trained model was applied, there were grip patterns from two collection

Table 5.1: Verification results in equal-error rate (%).

Train session	1+2	1+3	2+3
Test session	3	2	1
RF	13.1	5.5	3.2
High-pass filter	15.9	8.7	6.6
LABP, t_{opt} by method1	6.8	4.4	3.0
LABP, t_{opt} by method2	4.9	3.6	2.0

sessions in the training set. The value of t_{opt} was determined in three steps. First, one optimal value of t was decided with the data from one collection session for training, and those from the other collection session for testing. Second, another optimal value of t was decided in the same way, but with the data from these two sessions exchanged. Finally, t_{opt} was taken as the averaged value of the optimal t obtained in both steps. Therefore, unlike the first method, the second method took into account the grip-pattern variations across-sessions. That is, the across-session variations of grip patterns were modelled in the searching procedure of the optimal threshold with the second method.

The experimental results are presented in Table 5.1, with the grip-pattern images preprocessed by LABP and the high-pass filter respectively. As a reference the verification results, based on data not preprocessed by either of the methods, are listed as well, denoted as 'RF'. One can see that compared to the reference results, a significant improvement was achieved with the application of LABP, whereas the results were degraded using high-pass filtering. We believe that the substantial improvement of the verification performance based on LABP was mainly attributed to the two effects of this method, i.e. reducing the variation of pressure distribution across sessions and extracting information of the hand shape, as discussed in Section 5.2. Also, one can see that better verification results were obtained when the second method was used to find out t_{opt}, than those based on the first method. The reason was that the across-session variations of images were better modelled in the searching procedure with the second method than with the first method.

Further analysis shows that after LABP, the false-rejection rate of the

system decreased greatly, whereas its false-acceptance rate did not change very much. Figure 5.8 shows an example, where the grip patterns from the first and second collection sessions are used for training, and those from the third collection session for testing. Here the optimal value of t as 0.005 based on the second method is applied.

5.5 Conclusions

We proposed a novel technique as preprocessing for grip-pattern classification, namely, Local Absolute Binary Patterns (LABP). With respect to a certain pixel in an image, LABP processing quantifies how its neighboring pixels fluctuate. LABP has two main effects on the grip-pattern images. First, the across-session variation of pressure distribution is greatly reduced. The reason is that in each image after LABP the pressure values in different subareas within the hand part become much more equalized, compared to the case in the original image. Second, with LABP processing the information of hand shapes is extracted by enhancing the contrast between the grip patterns and the background. After LABP the grip pattern becomes much more distinctive with respect to the background. Attributed to these two effects, a significant improvement of the verification performance has been achieved.

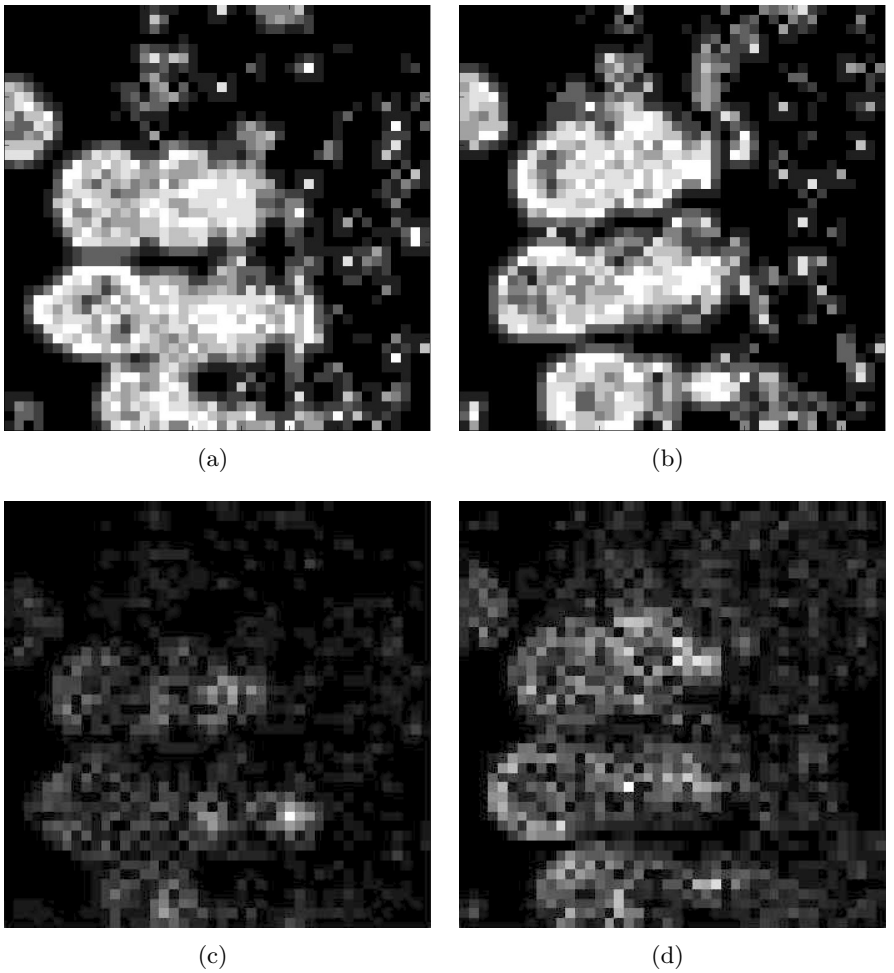

(a) (b)

(c) (d)

Figure 5.4: (a) Grip-pattern image in Figure 5.1(a) processed by LABP with $t = 0.006$. (b) Grip-pattern image in Figure 5.1(b) processed by LABP with $t = 0.006$. (c) Grip-pattern image in Figure 5.1(a) processed by high-pass filter with $\sigma = 0.5$. (d) Grip-pattern image in Figure 5.1(b) processed by high-pass filter with $\sigma = 0.5$.

(a)

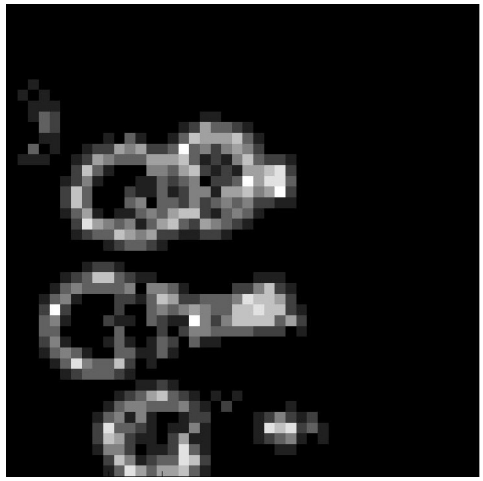

(b)

Figure 5.5: (a) Grip-pattern image in Figure 5.1(a) after LABP with $t = 0$. (b) Grip-pattern image in Figure 5.1(a) after LABP with $t = 0.02$.

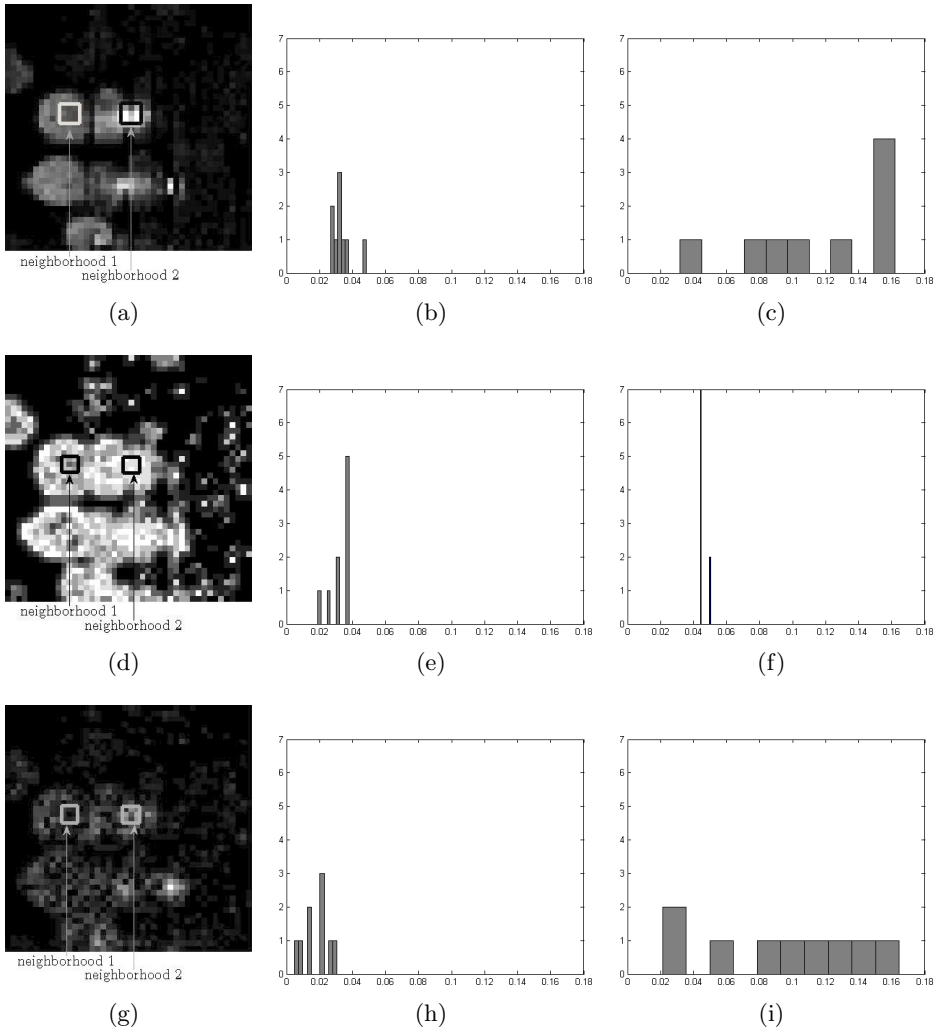

Figure 5.6: (a) Three-by-three neighborhood 1 and 2 in Figure 5.1(a). (b) Histogram of pixel values in neighborhood 1 in (a). (c) Histogram of pixel values in neighborhood 2 in (a). (d) Three-by-three neighborhood 1 and 2 in Figure 5.4(a). (e) Histogram of pixel values in neighborhood 1 in (d). (f) Histogram of pixel values in neighborhood 2 in (d). (g) Three-by-three neighborhood 1 and 2 in Figure 5.4(c). (h) Histogram of pixel values in neighborhood 1 in (g). (i) Histogram of pixel values in neighborhood 2 in (g).

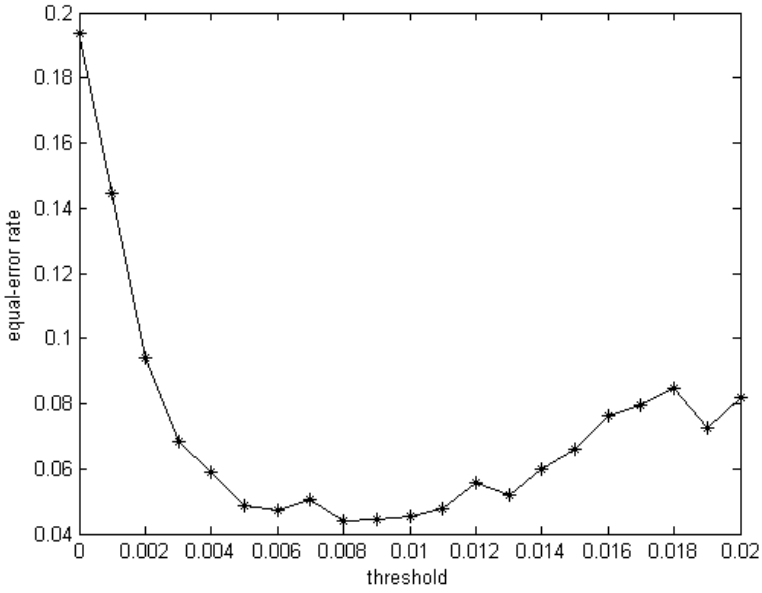

Figure 5.7: Equal-error rate obtained with different values of t

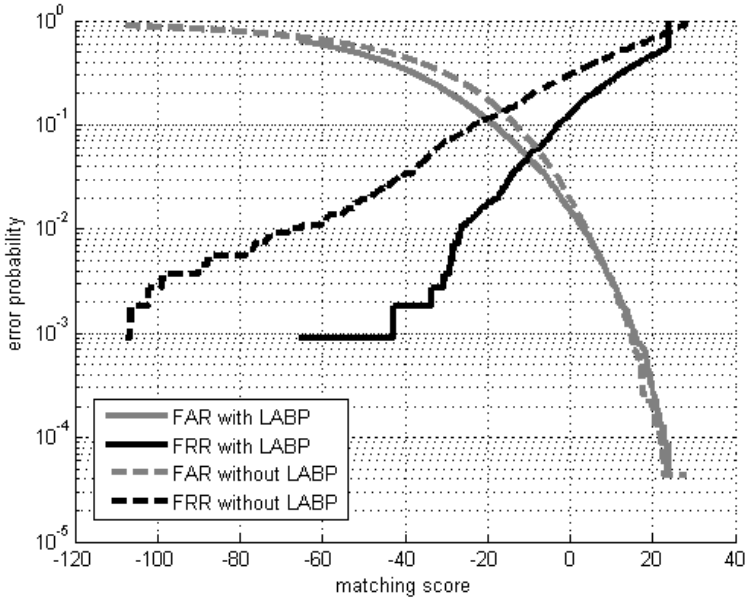

Figure 5.8: False-acceptance rate and false-rejection rate curves obtained with and without LABP, where $t = 0.005$.

Chapter 6

Grip-pattern recognition based on maximum-pairwise and mean-template comparison[1]

Abstract. *In this chapter grip-pattern verification is done using two types of comparison methods. One is mean-template comparison, where the matching score between a test image and a subject is computed by comparing this test image to the mean value of the subject's training samples. The other method is maximum-pairwise comparison, where the matching score between a test image and a subject is selected as the maximum of all the similarity scores, resulting from the comparisons between this test image and each training sample of the subject. Experimental results show that a much lower false-acceptance rate at the required false-rejection rate of the system, i.e.* 10^{-4}*, can be obtained using maximum-pairwise comparison, than mean-template comparison.*

[1]This chapter is based on [28].

6.1 Introduction

As described in Section 1.2, our verification system as part of a smart gun should have a very low false-rejection rate, rendering it highly unlikely that a police officer could not fire his or her own gun. The current official requirement in the Netherlands, for example, is that the probability of failure of a police gun be lower than 10^{-4}. Therefore, in our work the false-rejection rate for verification must remain below this value. Under this precondition, the false-acceptance rate should be minimized.

The experimental results for verification in previous chapters, however, have all been presented in terms of the equal-error rate, instead of the false-acceptance rate at the false-rejection rate equal to 10^{-4}. This is because in earlier chapters we mainly focused on improving the verification performance of the system in general, and we found that there was no conflict between it and reducing the false-acceptance rate at the false-rejection rate equal to 10^{-4}. That is, in our work a lower equal-error rate corresponded to a lower false-acceptance rate at the required false-rejection rate, and vice versa. Also, note that the experimental results for verification obtained so far, were all based on mean-template comparison. That is, the matching score between a test image and a subject was obtained by comparing this test image to the mean value of the subject's training samples.

In this chapter we propose another method of comparison: maximum-pairwise comparison. Using this method, a test grip-pattern image is compared to a number of training samples of a subject, one by one. Of all the similarity scores obtained, the highest one is selected as the matching score between this test image and the subject. We came up with this method since we expected better verification performance to be achieved, with the similarity between a test image and each *individual* training sample of a subject examined. This is because the mean value of a subject's training samples may not be able to represent the test data of this subject well, due to the model drift problem of grip patterns, as described in Chapter 3. The major advantage of maximum-pairwise comparison over mean-template comparison has proved to be, that a significantly lower false-acceptance rate for grip-pattern verification at the false-rejection rate equal to 10^{-4} can be achieved, even though the equal-error rates obtained with these two methods do not differ much.

This chapter analyzes the verification performance of grip-pattern recognition, using maximum-pairwise comparison and mean-template comparison, respectively. The remainder of this chapter is organized as follows. The verification algorithm, in both cases of mean-template comparison and maximum-pairwise comparison, will be briefly described in Section 6.2. Section 6.3 presents and discusses the experimental results. Finally, conclusions will be given in Section 6.4.

6.2 Verification algorithms

Mean-template comparison

If mean-template comparison is used, the verification algorithm is the same as was described in Chapter 3. In order to highlight the difference between mean-template comparison and maximum-pairwise comparison, we briefly repeat the verification algorithm as follows. It is assumed that the data are Gaussian and verification is based on a likelihood-ratio classifier. The likelihood-ratio classifier is optimal in the Neyman-Pearson sense, i.e. the false-acceptance rate is minimal at a given false-rejection rate or vice versa, if the data have a known probability density function [29], [30].

A grip-pattern image is arranged into a column vector. The matching score between a measurement $\mathbf{z}$ and a class c is derived from the log-likelihood ratio, computed by

$$\begin{aligned} M(\mathbf{z}, \mu_{\mathrm{c}}) &= -\frac{1}{2}(\mathbf{z}-\mu_{\mathrm{c}})' \boldsymbol{\Sigma}_{\mathrm{c}}^{-1}(\mathbf{z}-\mu_{\mathrm{c}}) \\ &\quad + \frac{1}{2}(\mathbf{z}-\mu_{\mathrm{T}})' \boldsymbol{\Sigma}_{\mathrm{T}}^{-1}(\mathbf{z}-\mu_{\mathrm{T}}) \\ &\quad - \frac{1}{2}\log|\boldsymbol{\Sigma}_{\mathrm{c}}| + \frac{1}{2}\log|\boldsymbol{\Sigma}_{\mathrm{T}}|, \end{aligned} \tag{6.1}$$

where μ_{c} and $\boldsymbol{\Sigma}_{\mathrm{c}}$ denote the mean vector and covariance matrix of class c; while the data of impostors are characterized by the total mean μ_{T} and total covariance matrix $\boldsymbol{\Sigma}_{\mathrm{T}}$ of grip patterns. The $'$ denotes vector or matrix transposition. If $M(\mathbf{z})$ is above a preset threshold, the measurement is accepted as being from the genuine user. Otherwise it is rejected [21].

Due to the large dimensionality of feature vectors (1936-dimensional vectors), a small-sample-size problem should be solved prior to classification [21]. We therefore apply the principal component analysis and linear discriminant analysis subsequently, to reduce the dimensionality [39]. This procedure can be represented by a transformation matrix $\mathbf{F}$. The total covariance matrix becomes an identity matrix after transformation, while the within-class covariance matrix diagonal. The feature vectors now have a dimensionality of $N_{\text{user}} - 1$, with N_{user} the number of subjects in the training set [21]. Thus, (6.1) becomes

$$\begin{aligned} M(\hat{\mathbf{z}}, \hat{\mu}_{\text{c}}) &= -\frac{1}{2}(\hat{\mathbf{z}} - \hat{\mu}_{\text{c}})' \mathbf{\Lambda}_{\text{R}}^{-1} (\hat{\mathbf{z}} - \hat{\mu}_{\text{c}}) \\ &\quad + \frac{1}{2}(\hat{\mathbf{z}} - \hat{\mu}_{\text{T}})' (\hat{\mathbf{z}} - \hat{\mu}_{\text{T}}) \\ &\quad - \frac{1}{2} \log |\mathbf{\Lambda_R}|, \end{aligned} \tag{6.2}$$

where

$$\hat{\mathbf{z}} = \mathbf{F}\mathbf{z}, \tag{6.3}$$

$$\hat{\mu}_{\text{c}} = \mathbf{F}\mu_{\text{c}}, \tag{6.4}$$

$$\hat{\mu}_{\text{T}} = \mathbf{F}\mu_{\text{T}}, \tag{6.5}$$

and $\mathbf{\Lambda}_{\text{R}}$ denotes the resulting diagonal within-class covariance matrix, which has been assumed to be class independent [21].

Maximum-pairwise comparison

If maximum-pairwise comparison is applied, the verification algorithm remains almost the same as in the case of mean-template comparison. We only need to modify (6.4), the expression of $\hat{\mu}_{\text{c}}$, to:

$$\hat{\mu}_{\text{c}} = \mathbf{F}\mathbf{s}_i, \tag{6.6}$$

where $\mathbf{s}_i, i = 1, ..., l$ are a number of training samples of class c. The matching score between a measurement $\mathbf{z}$ and a class c is expressed as

$$\max_i M(\hat{\mathbf{z}}, \mathbf{F}\mathbf{s}_i). \tag{6.7}$$

Table 6.1: Verification results with double-trained model applied.

Train session	Test session	Equal-error rate (%)		FAR*ref* (%)	
		MTC	MPWC	MTC	MPWC
2+3	1	2.0	2.6	63	10
1+3	2	3.6	3.3	65	50
1+2	3	4.9	4.4	60	50
Average		3.2	3.6	57	45

That is, a measurement $\mathbf{z}$ is compared to each of these training samples. Then, the highest similarity score is selected as the matching score between this measurement and class c.

6.3 Experiments, results and discussion

We did experiments based on the protocol described in Section 3.3. In each experiment both the false-acceptance rate at the false-rejection rate equal to 10^{-4}, and the equal-error rate for verification were measured. The local absolute binary patterns and template-matching registration were applied, prior to classification. In maximum-pairwise comparison, we compared a test image to all the training samples of a subject.

The experiments were done with and without the double-trained model applied, respectively. This was because of our expectation that the effects of maximum-pairwise comparison on the verification performance may differ in these two cases, since with double-trained model the number of training samples of a subject would become much greater and thus more comparisons would be made. Table 6.1 and 6.2 show the experimental results. 'FAR*ref*' represents the false-acceptance rate at the false-rejection rate equal to 10^{-4}. 'MTC' represents mean-template comparison, and 'MPWC' represents maximum-pairwise comparison. The 'Average' verification results were estimated based on the matching scores, obtained from all cases of the combinations of training and test sessions.

First, one can see that the major advantage of maximum-pairwise comparison over mean-template comparison is, that with it a much lower false-acceptance rate at the false-rejection rate equal to 10^{-4} can be achieved.

Table 6.2: Verification results without double-trained model applied.

Train session	Test session	Equal-error rate (%)		FAR*ref* (%)	
		MTC	MPWC	MTC	MPWC
2	1	2.1	2.4	65	48
3	1	6.6	5.8	70	55
1	2	4.7	4.1	85	82
3	2	7.1	7.6	55	55
1	3	7.1	7.3	65	50
2	3	10.0	9.5	70	70
Average		6.3	6.2	67	60

Figure 6.1 shows an example of the false-acceptance and false-rejection rate curves, using both comparison methods. Here, the grip patterns recorded in the second and third sessions are used for training, and those in the first session for testing. Note that since the data are not sufficient to reach the point where the false-rejection rate is equal to 10^{-4}, FAR*ref* has to be estimated by extrapolating the curves. Second, comparing the average results shown in Table 6.1 and 6.2, one can see that FAR*ref* decreased much more when the double-trained model was applied, than when it was not.

Further analysis indicated that the reason why FAR*ref* decreased was, that the matching scores of those images from the genuine users, which were of relatively low values with mean-template comparison applied, increased significantly with maximum-pairwise comparison applied instead. Figure 6.2, for example, shows the matching scores of grip-pattern images from the genuine users with maximum-pairwise comparison and mean-template comparison, respectively. Here, the grip patterns recorded in the first and second collection sessions are used for training, and those from the third session for testing. The histograms of these matching scores of images from the genuine users are shown in Figure 6.3.

One can see from Figure 6.2 and Figure 6.3, that in general the lower matching score a test image has with mean-template comparison, the more this score increases with maximum-pairwise comparison. This can be explained as follows. Suppose that a test image from a genuine user has a relatively low matching score with this user with mean-template compari-

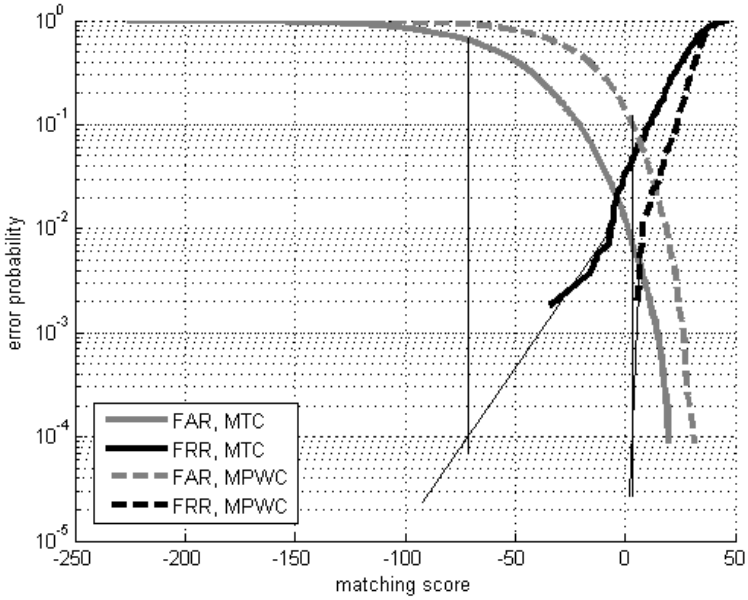

Figure 6.1: False-acceptance and false-rejection rate curves with both MPWC and MTC.

son, it must be due to a large mismatch between this test image and the mean value of the user's training samples. Thus, most likely there is a large mismatch between this image and most training samples of the user. However, as long as there exists at least one training sample that is fairly similar to the test image, the similarity score between the test image and this particular training sample can be of a relatively high value. Then, with maximum-pairwise comparison, this similarity score will be selected as the matching score between the test image and the genuine user. That is, the matching score between this image and the user will become relatively high.

Also, this can explain why FAR*ref* decreased more with the double-trained model applied, compared to the case where it was not applied. When two sessions of grip patterns were combined for training, the number of training samples became much larger. As a result, the probability that such a training sample existed that was similar to the test image, also

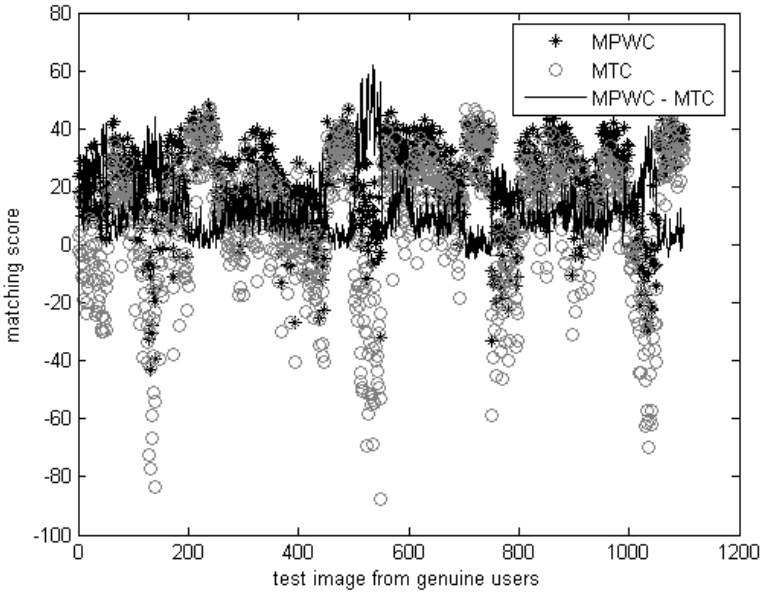

Figure 6.2: Matching scores of images from genuine users with both MPWC and MTC.

became larger.

The explanations given above can be demonstrated with an example in Figure 6.4. Here the matching score, between a test image from a genuine user and this user, is much higher with maximum-pairwise comparison than with mean-template comparison. Comparing all the images, one can see that there is a rather big "blob" in the middle-left part of both the mean image of the user's training samples, and the training sample that produces the minimum matching score, as shown in Figure 6.4(b) and Figure 6.4(d). Yet, in both the test image itself and the training sample that produces the maximum matching score, such a "blob" does not exist at the similar location, as shown in Figure 6.4(a) and Figure 6.4(c). The presence of the "blob" results from the fact, that during data recording a part of his palm near the wrist, sometimes, touched and exerted pressure onto the sensor wrapping around the grip of the gun. We believe that the absence of

the "blobs", among other factors, contributes to a relatively high matching score between the test image in Figure 6.4(a), and the image of the training sample shown in Figure 6.4(c).

Note that since FAR*ref* was estimated by extrapolation of the false-acceptance and false-rejection curves, its exact value may not be accurately measured. However, based on the observation and analysis we still believe that the maximum-pairwise comparison should result in a lower value of FAR*ref*, than the mean-template comparison.

6.4 Conclusions

Experiments for grip-pattern verification has been done with both mean-template comparison and maximum-pairwise comparison. A much lower false-acceptance rate was obtained with maximum-pairwise comparison, than with mean-template comparison, at the required false-rejection rate of our system, i.e. 10^{-4}.

This was attributed to the different working principles of these comparison methods. Even though a test image from a genuine user may have a large mismatch with the mean of this user's training samples, it is very likely that the similarity score between the test image and at least one training sample of this user is relatively high. Therefore, the computed matching score, between the test image and this user, is in general higher with maximum-pairwise comparison, than with mean-template comparison. And this, in turn, result in a lower false-acceptance rate at the aimed false-rejection rate.

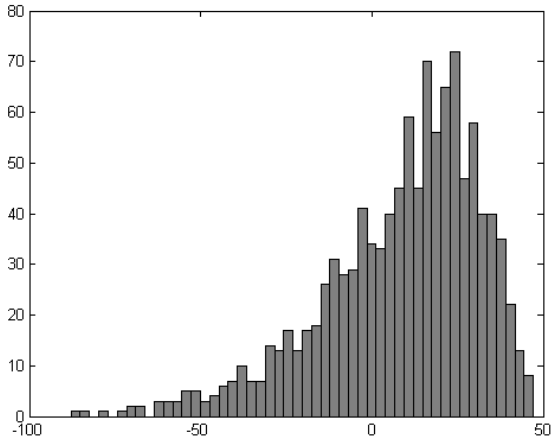

(a)

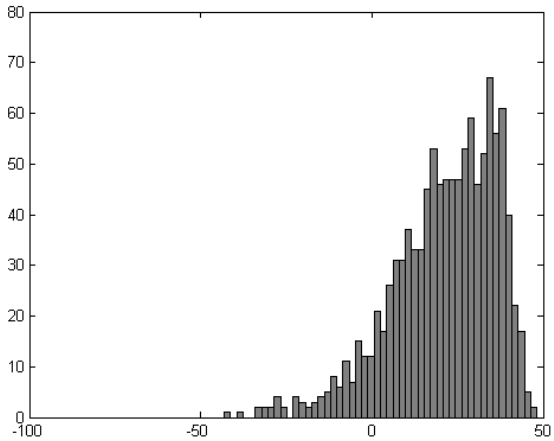

(b)

Figure 6.3: (a) Histogram of matching scores of images from the genuine users with MTC. (b) Histogram of matching scores of images from the genuine users with MPWC.

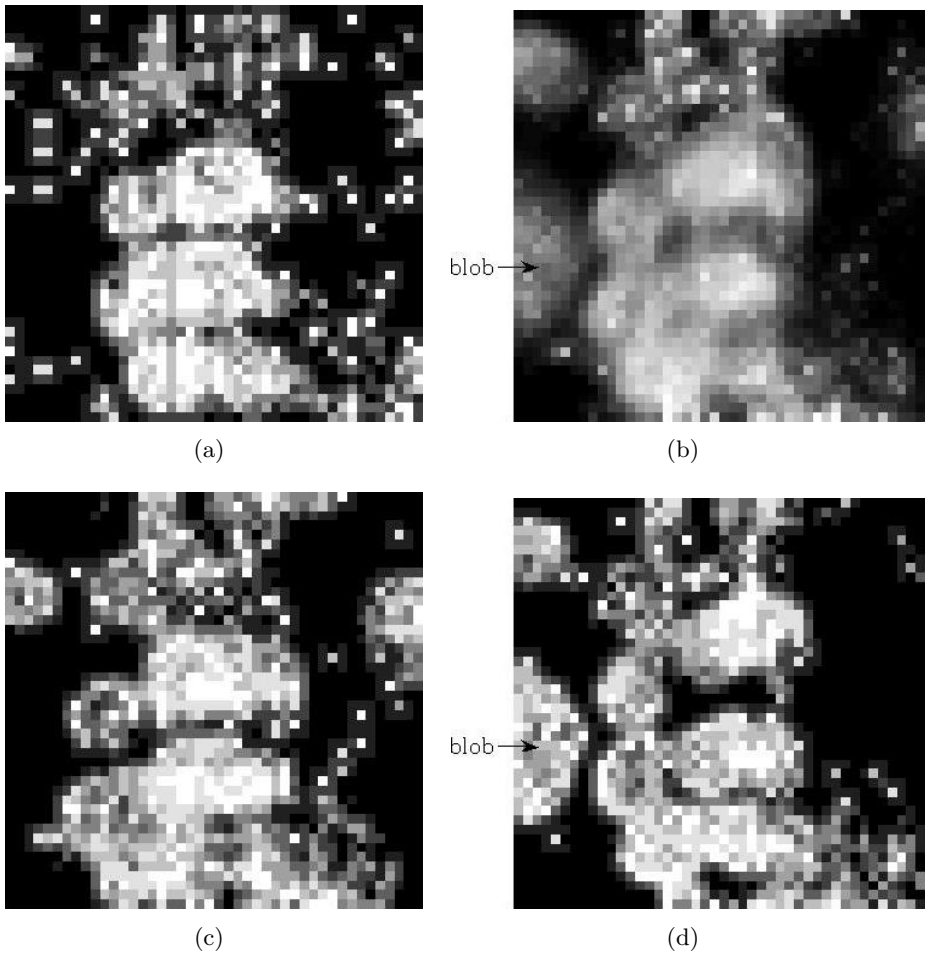

Figure 6.4: (a) Test image from a genuine user. (b) Mean of training samples. (c) Training sample producing the maximum matching score. (d) Training sample producing the minimum matching score.

Chapter 7

Restoration of missing lines in grip-pattern images[1]

Abstract. *A grip pattern is measured by a sensor sheet covering the grip of a gun. In practice missing lines can occur in a measured image, caused by various factors related to hardware implementation. This would degrade the verification performance of grip-pattern recognition. Therefore, restoration of the missing lines is necessary. In this chapter we propose a restoration method, based on minimization of the restoration error in the null space of data. Experimental results show that this method can effectively restore the missing lines in grip-pattern images. As a result, the verification performance of grip-pattern recognition can be significantly improved.*

7.1 Introduction

In data collection sessions, we observed that occasionally horizontal or vertical lines happened to be missing in a grip-pattern image. That is, the pressure values in those lines were all recorded as zeros. This was found to have been caused by some cable damage in the prototype, used to record the data. Figure 7.1, for example, shows a grip-pattern image where three horizontal lines are missing.

[1]This chapter is based on [23].

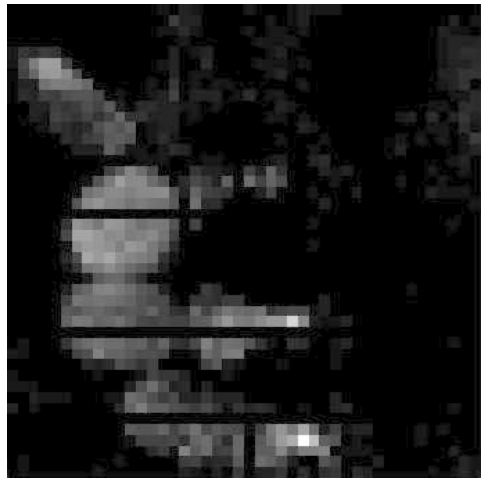

Figure 7.1: A grip-pattern image with three missing lines.

In practice missing lines can occur in a measured image, caused by various factors related to the hardware implementation in the gun. Just to name a few, consider defects of the measuring sensor, or a lose contact between two electric components embedded in the gun. The verification performance of grip-pattern recognition would then be degraded, due to the lack of data. Therefore, restoration of missing lines in a measured grip-pattern image is necessary to ensure correct operation in a practical situation.

In this chapter we propose a restoration method, based on minimization of the restoration error in the null space of the grip-pattern data. The idea has been inspired by the linear minimum variance estimators described in Chapter 2 of [53]. It requires training, based on a set of grip patterns without these missing lines. It also requires that the hardware provides a signal that indicates whether or not a line in a measured image is missing. Experimental results show that the missing lines in a grip-pattern image can be effectively restored, using this method. This has been proved in two ways. First, the average restoration error per pixel was much smaller, compared to the average within-class variance per pixel, estimated from data of a training set. Second, the verification performance of grip-pattern

recognition has been effectively improved, after restoration of the missing lines.

This chapter presents the restoration method and analyzes its effect on the verification performance of grip-pattern recognition. The remainder of this chapter is organized as follows. Section 7.2 describes the restoration algorithm. The experimental results will be presented and discussed in Section 7.3. Finally, conclusions will be given in Section 7.4.

7.2 Restoration algorithm

Let the column vector $\mathbf{x}$ denote a raw data vector with N elements. Furthermore, let

$$\mathbf{x} = \begin{pmatrix} \mathbf{u} \\ \mathbf{v} \end{pmatrix}, \tag{7.1}$$

be arranged such, that subvector $\mathbf{u}$ contains the m unknown data elements and subvector $\mathbf{v}$ the $N - m$ known ones. We assume that the covariance matrix of data, $\mathbf{\Sigma_{xx}}$, has a null space of a dimensionality of at least m, the number of unknown elements. Let the columns of $\mathbf{F} \in \mathbb{R}^N \times \mathbb{R}^M$ be an orthonormal basis of the observation space, e.g. obtained by performing PCA (Principal Component Analysis) on training data [39]. Let $\hat{\mathbf{u}}$ denote the estimate of $\mathbf{u}$. Our goal is to determine $\hat{\mathbf{u}}$, such that the null-space error, i.e. the norm of signal component in the null space after restoration

$$E_0 = \|(\mathbf{I} - \mathbf{F}\mathbf{F}^{'}) \begin{pmatrix} \hat{\mathbf{u}} \\ \mathbf{v} \end{pmatrix} \|^2 \tag{7.2}$$

is minimized. The symbol $'$ denotes vector or matrix transposition. The rationale behind this approach is that after restoration of the missing lines, an image will be as close as possible to the space spanned by $\mathbf{F}$ and, therefore, closely resembles a grip pattern.

The matrix $(\mathbf{I} - \mathbf{F}\mathbf{F}^{'})$ projects onto the null-space of $\mathbf{F}$, therefore we have

$$(\mathbf{I} - \mathbf{F}\mathbf{F}^{'})^2 = (\mathbf{I} - \mathbf{F}\mathbf{F}^{'}). \tag{7.3}$$

Now we partition $\mathbf{F}$ into two submatrices:

$$\mathbf{F} = \begin{pmatrix} \mathbf{F}_1 \\ \mathbf{F}_2 \end{pmatrix}, \tag{7.4}$$

with $\mathbf{F}_1$ an $m \times m$ matrix and $\mathbf{F}_2$ an $(N-m) \times m$ matrix. By using (7.3) and (7.4), the null-space error (7.2) can be rewritten as

$$\begin{aligned} E_0 &= (\hat{\mathbf{u}}' \mathbf{v}')(\mathbf{I} - \mathbf{F}\mathbf{F}') \begin{pmatrix} \hat{\mathbf{u}} \\ \mathbf{v} \end{pmatrix} && (7.5) \\ &= \|\hat{\mathbf{u}}\|^2 + \|\mathbf{v}\|^2 - \|\mathbf{F}_1' \hat{\mathbf{u}} + \mathbf{F}_2' \mathbf{v}\|^2. && (7.6) \end{aligned}$$

Setting the derivative of E_0 with respect to $\hat{\mathbf{u}}$ equal to zero, yields the following equation

$$(\mathbf{I} - \mathbf{F}_1 \mathbf{F}_1') \hat{\mathbf{u}} = \mathbf{F}_1 \mathbf{F}_2' \mathbf{v}, \tag{7.7}$$

from which $\hat{\mathbf{u}}$ can be solved if $(\mathbf{I} - \mathbf{F}_1 \mathbf{F}_1')$ has full rank.

Let us briefly look at the sensitivity of this restoration approach to the model errors. The model errors occur if the null-space components are present in the observed data. Assume that

$$\begin{aligned} \mathbf{x} &= \mathbf{x}_0 + \mathbf{g} && (7.8) \\ &= \begin{pmatrix} \mathbf{u}_0 \\ \mathbf{v}_0 \end{pmatrix} + \begin{pmatrix} \mathbf{g}_1 \\ \mathbf{g}_2 \end{pmatrix}, && (7.9) \end{aligned}$$

where $\mathbf{x}_0$ denotes the uncorrupted data and $\mathbf{g}$ an error component in the null space of $\mathbf{F}$. Therefore,

$$\mathbf{F}_1' \mathbf{g}_1 + \mathbf{F}_2' \mathbf{g}_2 = \mathbf{0}. \tag{7.10}$$

The error component after restoration becomes

$$\begin{aligned} \epsilon &= (\mathbf{I} - \mathbf{F}_1 \mathbf{F}_1')^{-1} \mathbf{F}_1 \mathbf{F}_2' \mathbf{g}_2 && (7.11) \\ &= -(\mathbf{I} - \mathbf{F}_1 \mathbf{F}_1')^{-1} \mathbf{F}_1 \mathbf{F}_1' \mathbf{g}_1 && (7.12) \\ &= \mathbf{g}_1 - (\mathbf{I} - \mathbf{F}_1 \mathbf{F}_1')^{-1} \mathbf{g}_1. && (7.13) \end{aligned}$$

The first term in (7.13) reflects the presence of model errors in the part that is to be restored. The second term reflects the error due to the compensation

for model errors in the given data. The compensation error grows strongly as the eigenvalues of $\mathbf{F}_1\mathbf{F}_1^{'}$ approach 1.

When using this algorithm for restoration of missing lines in grip-pattern images, we choose the training set such that it does not contain images with missing lines, and the missing lines are only present in images of the test set. The matrix $\mathbf{F}$ can, therefore, be derived from the training data.

7.3 Experiment, results and discussion

This section presents and discusses the experimental results, with the missing lines in grip-pattern images restored using our proposed method. The restoration quality will be evaluated by both the average restoration error per pixel, computed by

$$E = \frac{1}{m}\|\hat{\mathbf{u}}_i - \mathbf{u}_i\|^2, \tag{7.14}$$

and effect of the application of the restoration method on the verification performance of grip-pattern recognition. In our experiments, the grip-pattern data recorded in the first and second sessions were used for training and testing, respectively.

Before applying the restoration method, we first set the pressure values of a number of lines to be zeros in each image of the test set. That is, we deliberately made some lines in an image to be missing. And, in each image these lines were randomly selected. This was done in ten cases, where the pressure values in one up to ten lines were set to be zeros, respectively. Figure 7.2, as an example, illustrates the restoration effect of this method on a grip-pattern image, where four lines are deliberately made to be missing.

Table 7.1 and Figure 7.3 show the average restoration errors per pixel in all ten cases. As a reference the average within-class variance per pixel, σ_{W}^2, is also presented. It is estimated from data of the training set as

$$\sigma_{\mathrm{W}}^2 = \frac{1}{N \cdot N_{\mathrm{user}}} \sum_j \sum_k \|\mathbf{x}_{j,k} - \overline{\mathbf{x}}_j\|^2. \tag{7.15}$$

N_{user} denotes the number of samples in the training set. $\mathbf{x}_{j,k}$ and $\overline{\mathbf{x}}_j$ represent the k^{th} sample and the mean value of class j in the training set,

Table 7.1: Restoration error per pixel.

Number of missing lines	E
1	159.0
2	165.9
3	170.5
4	164.2
5	165.0
6	164.9
7	166.1
8	165.2
9	166.2
10	162.8
σ_{W}^2	253.8

respectively. One can see that compared to the average within-class variance per pixel, σ_{W}^2, the average restoration error per pixel was much smaller. This indicates that the missing data have been restored effectively, using our proposed method. One can also see, that the average restoration error per pixel was fairly constant, with the number of missing lines in an image varying from one up to ten. This suggests that our proposed restoration method is robust against the increase of the number of missing lines in an image.

In order to investigate the effect of the restoration method on the verification performance of grip-pattern recognition, we performed the following experiment. First, the grip-pattern verification was done on the test set, where a number of lines were randomly selected and deliberately made to be missing in each image. Second, the test set composed of grip-pattern images with restored lines was used for verification. Third, the verification performance obtained in the first and second steps was compared. This whole procedure was repeated in ten cases, with the number of missing lines in each test image increasing from one up to ten. The verification performance was evaluated by the overall equal-error rate of all the subjects. The experimental results are shown in Table 7.2. As a reference, the verification result on the test set, where the complete data are available,

Table 7.2: Verification results in equal-error rate (%)

Complete test set	7.9	
Number of missing lines	With missing lines	With restored lines
1	8.6	8.0
2	9.1	8.0
3	9.4	8.2
4	9.6	8.4
5	10.4	8.5
6	10.9	8.7
7	11.7	8.8
8	12.6	9.0
9	12.8	9.1
10	13.5	9.3

is also presented. The verification performance as the number of missing lines in each image increases is further illustrated in Figure 7.4.

Table 7.2 and Figure 7.4 show, that verification performance of grip-pattern recognition has been significantly improved, with the missing lines in a measured image restored using the proposed method. And, the relative improvement became more as the number of missing lines increased. Therefore, the restoration method is especially helpful with a relatively large number of missing lines in an image.

Also, one can see that the verification performance with the restored data became worse, as the number of missing lines in each image increased. The is because the more lines were restored in each image, the higher the total restoration error there was per image. However, one can also see that the degradation was not much. This suggests again, that our proposed restoration method is robust against the increase of the number of missing lines in an image.

7.4 Conclusions

Since in practice missing lines can occur in a measured image for grip-pattern verification, caused by various factors related to hardware imple-

mentation, restoration of these lines is necessary to ensure correct operation in a practical situation. Therefore, we have proposed a restoration method, based on minimization of the restoration error in the null space of data.

Experimental results have shown that the proposed method can effectively restore the missing lines in a grip-pattern image. This has been reflected in two ways. First, the average restoration error per pixel was much smaller, compared to the average within-class variance per pixel. Second, the verification performance of grip-pattern recognition has been improved significantly, with the missing lines in an image restored using the proposed method. We have also found that the restoration method is fairly robust, against the increase of the number of missing lines in an image.

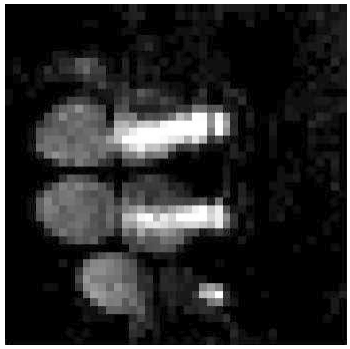

(a)

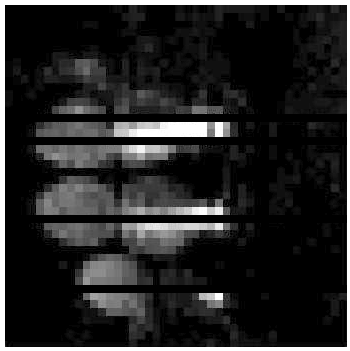

(b)

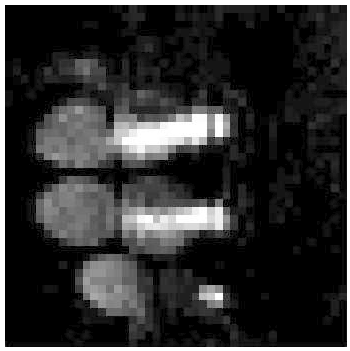

(c)

Figure 7.2: (a) A grip-pattern image with no missing lines. (b) The grip-pattern image with four lines set to be zeros. (c) The grip-pattern image with missing lines restored.

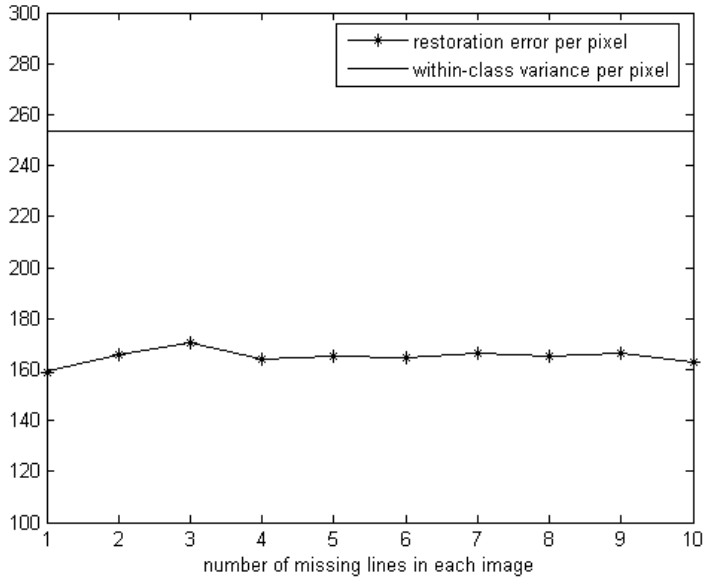

Figure 7.3: Average restoration error per pixel and within-class variance per pixel.

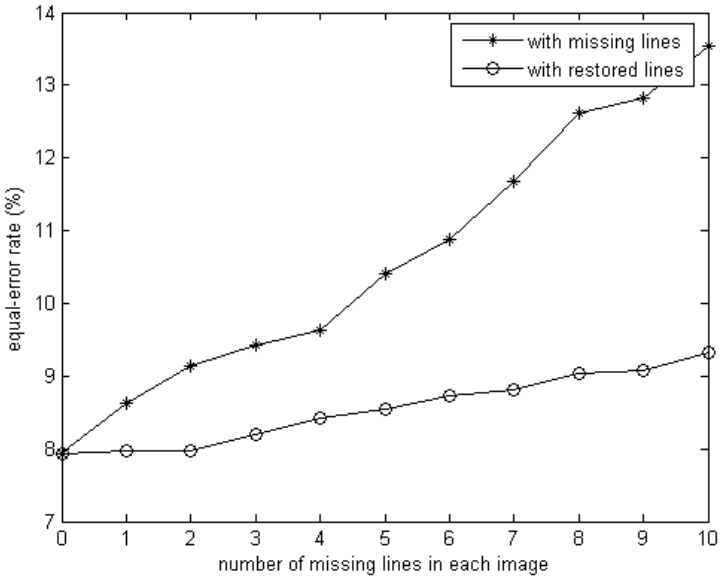

Figure 7.4: Verification performance using test data with missing and restored lines, respectively.

Chapter 8

Grip-pattern verification based on high-level features

Abstract. *In this chapter we base the grip-pattern verification on high-level features, namely, physical characteristics of the hand or hand pressure extracted from a grip-pattern image. Specifically, two types of high-level features are investigated. First, certain features extracted from the correlation image, between a test image and a part of the grip pattern in an image as a template. Second, the contour of the fingers in a grip-pattern image. Experimental results show, that the verification performance based on these high-level features is generally not good. This suggests, that the high-level features extracted from a grip-pattern image are not discriminative enough for identity verification.*

8.1 Introduction

So far in our work the grip-pattern verification was all based on the low-level features, i.e. the pressure values measured as the output of the sensor. Here, in this chapter we investigate the verification performance of the system by using the high-level features extracted from the grip-pattern images. By "high-level features" we refer to the physical characteristics of the hand or hand pressure of a subject obtained from a grip pattern image.

Particularly, two types of high-level features will be investigated in this chapter. First, the verification is done based on the features extracted from the correlation image, between a test image and a part of the grip pattern in an image as a template. This is mainly based on our expectation that the local pressure distribution within a small subarea of an image does not vary much across sessions, even though a great across-session variation of the pressure distribution in a whole image occurs. The second type of high-level feature is the contour of the fingers in a grip-pattern image. This is mainly inspired by our observation that the hand shape of a subject remains constant across sessions. Also, the feasibility of hand geometry recognition based on the contour of a hand has been demonstrated, for example, in [16].

In the remainder of this chapter we will present and discuss the experimental results of the grip-pattern verification by using these two types of high-level features, respectively. Based on the analysis, conclusions will be given afterwards.

8.2 Grip-pattern verification based on correlation features

Verification was done based on the features extracted from the correlation image between a test image and a part of a image as a template. Specifically, besides the rectangle that contained the pattern of the thumb, each rectangle which contained the pattern of the tip, middle, or lower section of a finger on the mean image of the training samples of a subject, respectively, was considered as a template rectangle of this subject.

Let $\mathbf{x}$ and $\mathbf{y}$ denote a test image and a template rectangle, respectively. Computation of the correlation image between them can be described as follows. First, define the correlation result $\mathbf{z}$ as an image of the same size as $\mathbf{x}$. Second, put $\mathbf{y}$ on top of $\mathbf{x}$ in a certain position. Multiply the value of each pixel in $\mathbf{y}$ by the corresponding value of the pixel in $\mathbf{x}$ underneath. Let $\mathbf{z}_i$ denote the pixel in $\mathbf{z}$, located at the same position as the pixel in $\mathbf{x}$ underneath the center pixel of $\mathbf{y}$. The value of $\mathbf{z}_i$ is then computed as

$$\mathbf{z}_i = \frac{\mathbf{y}^{'} \cdot \mathbf{x}_{i,\text{underneath}}}{\|\mathbf{y}\|\|\mathbf{x}_{i,\text{underneath}}\|}, \tag{8.1}$$

where the $'$ denotes vector or matrix transposition, and $\mathbf{x}_{i,\text{underneath}}$ denotes the rectangular subarea of $\mathbf{x}$ underneath $\mathbf{y}$. Finally, slide $\mathbf{y}$ along $\mathbf{x}$ and compute every other pixel in $\mathbf{z}$ in the same way. As an example, Figure 8.1 shows a template rectangle, a test image and their correlation image.

Subsequently, different features of the correlation image were measured. These features were: the maximum value in the whole image, the maximum value in a defined scope of the image, the position of the maximum value in the image, and the position of the maximum value in a defined scope of the image. With the "Defined scope" we refer to an area in the correlation image that either exactly overlaps the template rectangle, or a subarea of it with the same number of pixels shrunk from each side.

Finally, based on each of the four extracted features mentioned above, two types of classifiers were implemented for verification respectively. One was the correlation classifier, where the matching score between a test image and a template was the value of a certain type of feature measured in their correlation image. The other one was the likelihood-ratio classifier, where the likelihood ratio of a test image and a template was computed as the matching score [21]. The verification results based on these two types of classifiers are discussed as follows.

Correlation classifier based

Based on the correlation classifier the experiment was done, with the data collected in the first and second sessions used for training and testing, respectively. The verification performance was evaluated by the overall equal-error rate of all the subjects. It was computed from the values of the features measured in the correlation images of all the genuine users and impostors. The template rectangles were manually selected.

We found that no matter which type of feature was measured, the best verification results were those based on the correlation between a test image and the pattern of the tip section of the middle, ring, or little finger, compared to the case where the pattern of the middle or the lower section of a finger was used as the template image. This is easy to understand,

Table 8.1: Best verification results in equal-error rate (%) based on correlation classifier.

	Feature 1	Feature 2
Mtip as template	36.5	26.1
Rtip as template	26.8	26.6
Ltip as template	28.5	25.7

because these three parts of the grip-pattern image include more discriminative information than the others. We also found that the verification results were better if the maximum value in a defined scope of the correlation image, or the position of the maximum value in the correlation image was measured, compared to the case where some other type of feature was measured. To summarize, Table 8.1 shows the verification results which were relatively better, than those in other cases with an equal-error rate generally greater than 40%. 'Feature 1' represents the position of the pixel with the maximum value in the correlation image, and 'Feature 2' represents the maximum value in a defined scope measured in the correlation image. The tip sections of the middle, ring and little fingers are denoted by 'Mtip', 'Rtip' and 'Ltip', respectively.

One should also note that the verification results varied a lot with different number of pixels shrunk from each side of the template rectangle, in the case where the maximum value in a defined scope of the correlation image was measured. It was shown that the smaller the defined scope was, the better the verification result became. Let s denote the number of pixels shrunk from each side of the template rectangle. Table 8.2, as an example, shows the equal-error rate at different values of s based on the correlation between a test image and the pattern of the tip section of the ring finger as a template. 'NA' represents the case where the maximum value was searched on the whole correlation image. One can see that as the value of s became greater, the verification result became better.

Further analysis showed that the pixel of the maximum value in the whole correlation image was mostly located near the center of the template rectangle, if the test image came from a genuine user; while it was generally located outside the area of the template, if the test image came from an imposter. See Figure 8.2 for an example, with the pattern of the tip section

Table 8.2: Verification results at different values of s, with maximum value measured in a defined scope of correlation image.

s	NA	0	1	2	3
Equal-error rate (%)	39.9	36.2	33.7	30.1	26.6

of the ring finger in Figure 8.1 as the template. Here the same image as used in Figure 8.1 is selected as the image from the genuine user. Therefore, as the defined scope became smaller, the probability became lower that an image from an imposter was misclassified as from the genuine user; while it became higher that an image from the genuine user was misclassified as from an imposter. That is, when s became greater the false-acceptance rate of the verification decreased, while the false-rejection rate increased. However, the effect of the decrease proved to be stronger than that of the increase, as shown in Figure 8.3. This explains the improvement of the verification result as the value of s increased in Table 8.2.

One can see from Table 8.1 that the best verification result based on the correlation classifier, is of an equal-error rate equal to about 28% on average. Obviously it is not good. This suggests that our high-level features extracted are not discriminative enough for identity verification, compared to the low-level features. We also did the experiment for verification by combining the results in Table 8.1 using the fusing method [46], [48]. Unfortunately, this only reduced the equal-error rate to about 22%.

Likelihood-ratio classifier based

Verification was also done using the likelihood-ratio classifier, based on the features and templates which produced the relatively good results using the correlation classifier (see Table 8.1) [21]. Specifically, the input vector of the likelihood-ratio classifier was composed of: the maximum value in a defined scope of the correlation image between a test image and the patterns of the tip sections of the middle, ring, and little fingers, respectively; and the reference position (i.e. the x and y coordinates) of the maximum value in the correlation image between the test image and the patterns of the tip sections of the ring and little fingers, with respect to that in the correlation image between the test image and the pattern of the tip section

of the middle finger. Note that here we used the reference positions of the maximum values instead of their absolute positions, in order to reduce the effect of data variations across sessions caused by hand shift.

As earlier the data in the first and second collection sessions were used for training and testing, respectively. The verification performance was evaluated by the overall equal-error rate of all the subjects. It was computed from the likelihood ratios of all the genuine users and impostors. The data was transformed by PCA (Principal Component Analysis) and LDA (Linear Discriminant Analysis) in exactly the same way as described in Section 3.2. However, the experimental result was not good, with an equal-error rate of about 36%.

8.3 Grip-pattern verification based on finger contour

Besides the features measured in the correlation image as mentioned above, verification experiments were also done based on the contour of the fingers. This was inspired by our observation that the hand shape of a subject remains constant across sessions. Prior to the contour extraction, as a first step, we applied the local absolute binary patterns (LABP) preprocessing to the grip-pattern image, so that the contour could be better extracted from the image. Second, the contour of the finger patterns were extracted by the Canny edge detector. See Figure 8.4. Third, the landmarks were manually selected on the contour extracted. We selected 17 landmarks in each image: the finger tips of the middle, ring and little fingers; the connecting points of each adjacent two of these three fingers; and the other 12 points which were divided into three groups with each four as one group and equally distributed on both sides of a finger tip.

Verification was done using the likelihood-ratio classifier based on the x and y coordinates of the landmarks [21]. Still we used the data in the first and second collection sessions for training and testing, respectively. The verification performance was evaluated by the overall equal-error rate of all the subjects, computed from the likelihood ratios of all the genuine users and impostors. PCA and LDA transformations were made to the data in exactly the same way as described in Section 3.2.

However, the experimental result proved not good, with an equal-error rate of about 21%. As in the case where the verification was based on features measured in the correlation image, this suggests that our high-level features are not discriminative enough for identity verification

8.4 Conclusions

Grip-pattern verification based on two types of high-level features are investigated. One is the features measured in the correlation image between a test image and a part of the grip pattern in an image as a template. The second type of high-level feature is the contour of the finger patterns in an image. The experimental results show that the verification performance based on high-level features is generally not good. This is due to the fact that the high-level features extracted from the grip-pattern images are not discriminative enough for identity verification, compared to the low-level features.

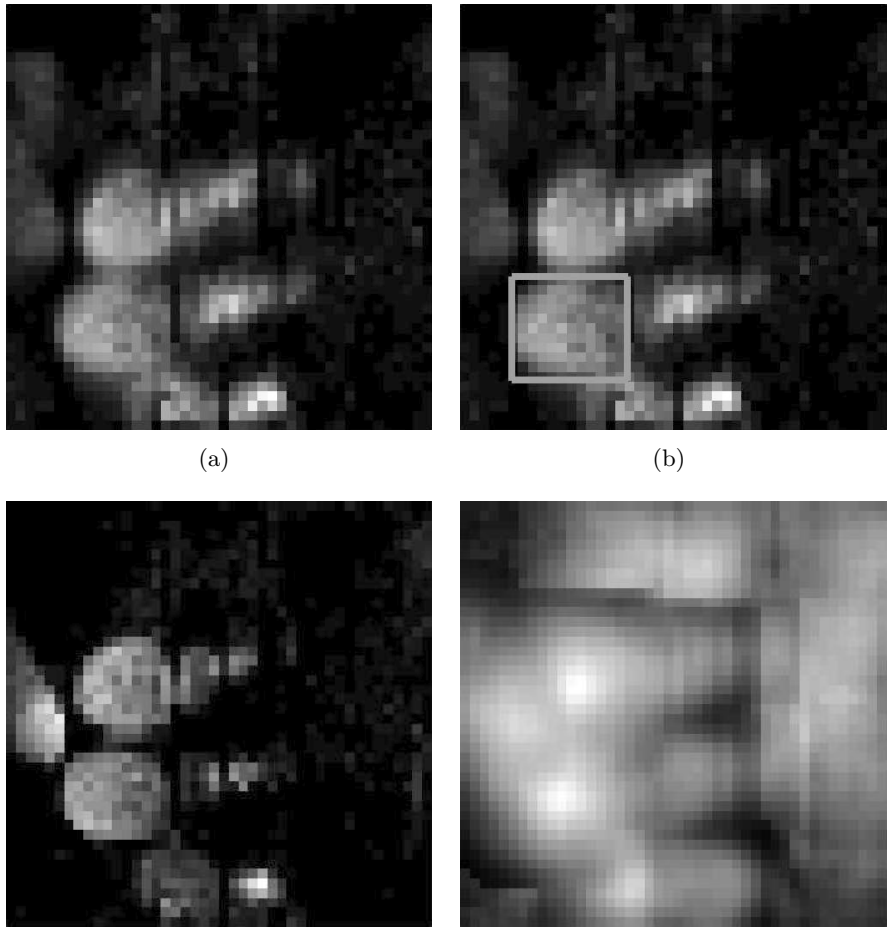

(a) (b)

(c) (d)

Figure 8.1: (a) Mean image of a subject from training session. (b) Tip section of ring finger on the mean image as template. (c) Image for testing. (d) Resulting image from correlation between the test image and the template.

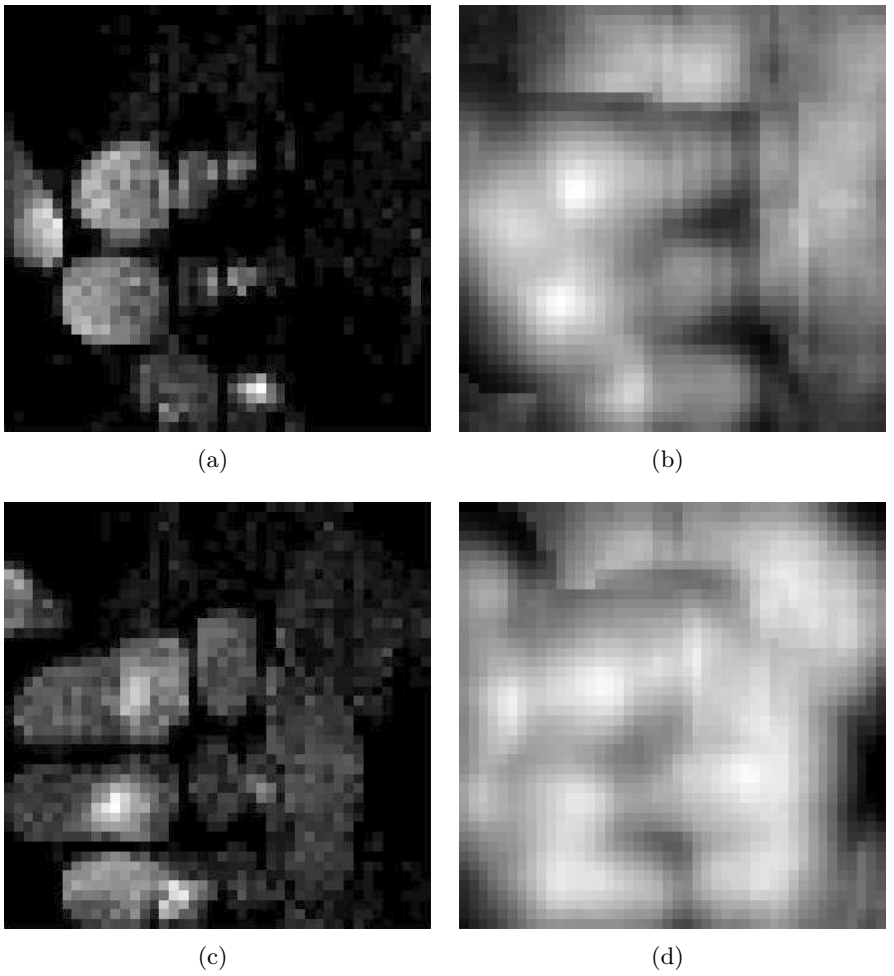

(a) (b)

(c) (d)

Figure 8.2: (a) Image from a genuine user. (b) Resulting image from correlation between the template and the image from the genuine user. (c) Image from an imposter. (d) Resulting image from correlation between the template and the image from the imposter.

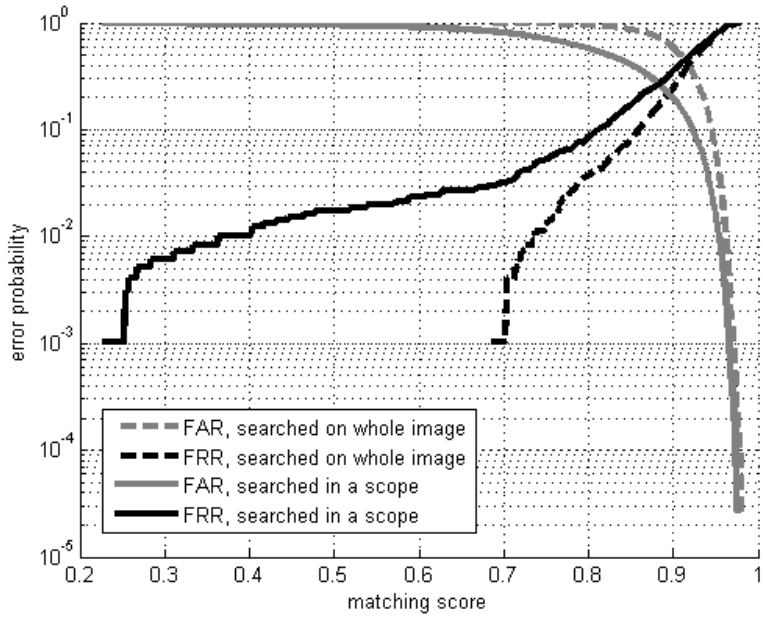

Figure 8.3: False-acceptance and false-rejection rate curves obtained based on the maximum value of the correlation image, searched on the whole image and within a defined scope at $s = 3$.

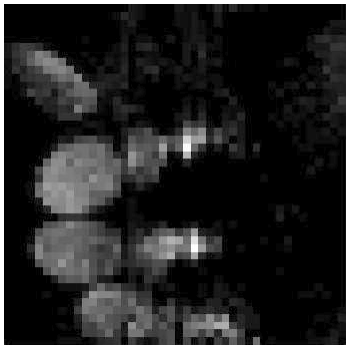

(a)

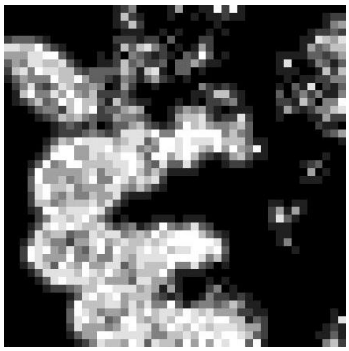

(b)

(c)

Figure 8.4: (a) An original grip-pattern image. (b) The image after LABP preprocessing. (c) Contour extracted by edge detector from the image after LABP preprocessing.

Chapter 9

Comparison of grip-pattern recognition using likelihood-ratio classifier and support vector machine[1]

Abstract. *In earlier work, grip-pattern verification was done by using a likelihood-ratio classifier. Since this type of classifier is based on estimation of the data probability density from a set of training samples, the verification results were not good due to the data drifting problem. Specifically, there occurred large variations of pressure distribution and hand position between the training and test grip-pattern images of a subject. In this chapter, grip-pattern verification using the support vector machine is done. Also, the verification performance is compared to that of the likelihood-ratio classifier. It has been shown that the support vector machine has better verification performance than the likelihood-ratio classifier, if there are considerable variations between the training and test images of a subject. However, once these variations are reduced or better modelled during training of the classifier by certain techniques, the support vector machine tends to lose its*

[1]This chapter is based on [27].

superiority.

9.1 Introduction

In Chapter 3, 4, 5, 6 and 7, the grip-pattern verification was done using a likelihood-ratio classifier. This type of classifier requires estimation of the probability density of grip patterns from a set of training samples. The verification performance therefore depends largely on how similar the probability density estimated from the training data is, to the actual distribution of the test data. That is, this type of classifier does not have good generalization performance and the verification performance will be degraded if data drifting occurs. Since there were large variations of pressure distribution and hand position between the training and test images of a subject, the verification performance of grip patterns was not good using the likelihood-ratio classifier. Figure 9.1 illustrates the grip-pattern variations across sessions.

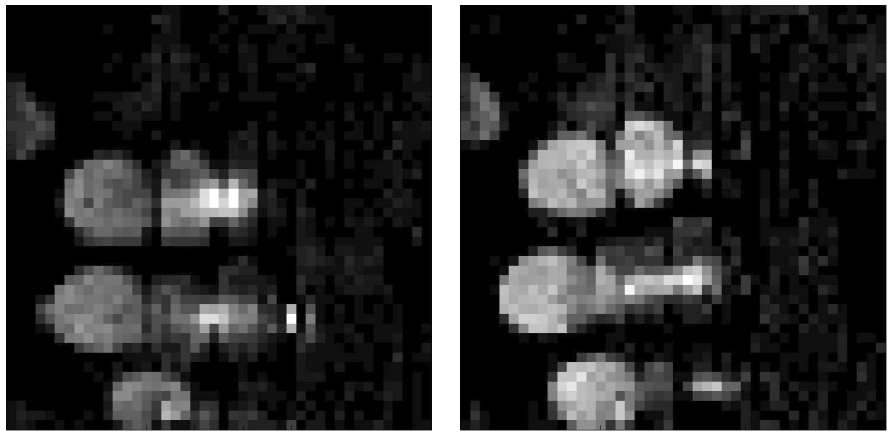

(a) (b)

Figure 9.1: Grip-pattern images of a subject in different sessions.

In order to investigate whether another type of classifier, based on a different principle, has better generalization performance, we implemented the

support vector machine for grip-pattern verification. The support vector machine, as a contrast to the likelihood-ratio classifier, does not estimate the data distribution. Instead, it tries to maximize the margin between different classes [32], [33], [54], [55], [56], [57]. It has been proved more capable of coping with the problem of data drifting, than other pattern-recognition classifiers [34], [35], [36], [37], [38].

In this chapter we compare the verification performance of grip-pattern recognition using the support vector machine and the likelihood-ratio classifier. The remainder of this chapter is organized as follows. Section 9.2 briefly describes the support vector machine and its application in the grip-pattern verification. Next, the experimental results will be presented and discussed in Section 9.3. Finally, conclusions will be given in Section 9.4.

9.2 Grip-pattern verification using support vector machine

Linear support vector machine for separable data

A support vector machine constructs a separating hyperplane, which maximizes the margin between two data sets [32], [33]. We start with the simplest case: linear support vector machine trained on separable data. We label a set of training data as $(\mathbf{x}_i, y_i), i = 1, ..., l.$ $\mathbf{x}_i \in \mathbf{R}^d$ and $y_i \in \{1, -1\}^l$ denote a d-dimensional data vector and its class label, respectively.

Suppose we have a hyperplane that separates the data of two classes, as shown in Figure 9.2. The data vectors $\mathbf{x}$ that lie on the hyperplane satisfy

$$\mathbf{w}^{'}\mathbf{x} + b = 0, \tag{9.1}$$

where $\mathbf{w}$ is the normal to the hyperplane, and $|b|/\|\mathbf{w}\|$ is the perpendicular distance from the hyperplane to the origin. The $'$ denotes vector or matrix transposition. Suppose that all the training data satisfy the following constraints:

$$\mathbf{w}^{'}\mathbf{x}_i + b \geq +1, \text{if } y_i = +1 \tag{9.2}$$

$$\mathbf{w}^{'}\mathbf{x}_i + b \leq -1, \text{if } y_i = -1. \tag{9.3}$$

Let d_+ and d_- be the shortest distances from the separating hyperplane to the closest samples of class 1 and -1, respectively. We define the "margin" of a separating hyperplane to be $d_+ + d_-$. In this linearly separable case, the support vector algorithm looks for the separating hyperplane with the largest margin between two classes.

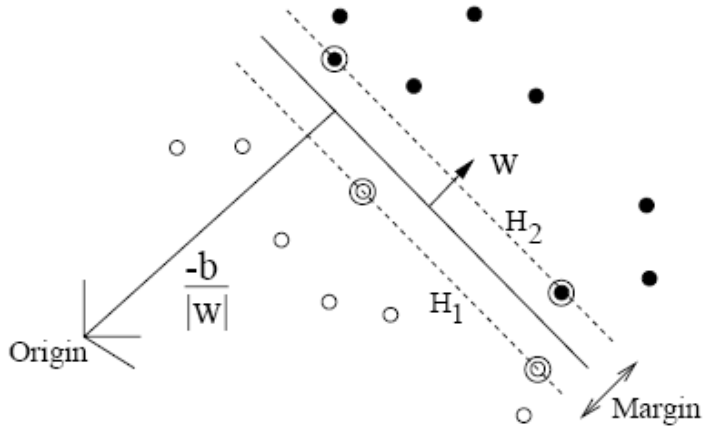

Figure 9.2: Linear separating hyperplanes for the separable case. The support vectors are circled.

Now consider the data vectors for which the equality in (9.2) holds. These vectors lie on the hyperplane $H_1 : \mathbf{w}'\mathbf{x} + b = 1$, with normal $\mathbf{w}$ and perpendicular distance from the origin $|1 - b|/\|\mathbf{w}\|$. The vectors, for which the equality in (9.3) holds, lie on the hyperplane $H_2 : \mathbf{w}'\mathbf{x} + b = -1$, also with normal $\mathbf{w}$ and perpendicular distance from the origin $| - 1 - b|/\|\mathbf{w}\|$. Therefore, $d_+ = d_- = 1/\|\mathbf{w}\|$, and the margin is $2/\|\mathbf{w}\|$. The pair of hyperplanes, which give the maximum margin, can thus be found by minimizing $\|\mathbf{w}\|^2$ subject to the constraints in (9.2) and (9.3). Note that H_1 and H_2 are parallel because they have the same normal, and that no data vectors fall between them.

After solving the constrained optimization problem described above, the linear separating hyperplane can be expressed as:

$$\mathbf{w}'\mathbf{x} + b = 0, \tag{9.4}$$

where

$$\mathbf{w} = \sum_{i=1}^{N_S} \alpha_i y_i \mathbf{s}_i. \tag{9.5}$$

Here, $\mathbf{s}_i$ denotes one of a set of N_S support vectors, which are the training vectors lying on H_1 and H_2; $\alpha_i > 0$ denotes its coefficient weight; and y_i denotes its class label [32], [33]. This hyperplane lies half way between H_1 and H_2 and parallel to them, as shown in Figure 9.2.

In the test phase given a test vector $\mathbf{x}$, we simply decide on which side of the decision boundary it lies, and assign $sgn(\mathbf{w}^{'}\mathbf{x} + b)$ as its class label.

Linear support vector machine for non-separable data

If the training data of two classes are non-separable, the decision boundary can be determined by adding a further cost to classification errors. A positive slack variable $\xi_i, i = 1, ..., l$ is introduced into the constraints, which then become:

$$\begin{aligned} \mathbf{w}^{'}\mathbf{x}_i + b &\geq +1 - \xi_i, \text{if } y_i = +1 && (9.6)\\ \mathbf{w}^{'}\mathbf{x}_i + b &\leq -1 + \xi_i, \text{if } y_i = -1 && (9.7)\\ \xi_i &\geq 0, \forall i. && (9.8) \end{aligned}$$

Therefore, for an error to occur, the corresponding ξ_i must exceed unity. This makes $\sum_i \xi_i$ the upper bound on the number of training errors. An extra cost for errors can be assigned by changing the objective function to be minimized to be

$$\frac{1}{2}\|\mathbf{w}\|^2 + C(\sum_i \xi_i)^k, \tag{9.9}$$

where C is a parameter to be chosen by the user. A larger C corresponds to assigning a higher penalty to errors.

Again, the linear separating hyperplane can be expressed as:

$$\mathbf{w}^{'}\mathbf{x} + b = 0, \tag{9.10}$$

where

$$\mathbf{w} = \sum_{i=1}^{N_S} \alpha_i y_i \mathbf{s}_i. \tag{9.11}$$

The only difference from the case of the optimal hyperplane is that the α_i now has an upper bound of C [32], [33].

Nonlinear support vector machine

Let us further assume that the decision boundary separating the classes is nonlinear. We map the data to a higher (possibly infinite) dimensional Euclidean space $\mathbf{\Gamma}$, with function Φ. In this space, a linear separating hyperplane will be found, with the maximum margin between two classes. That is, a linear support vector machine will be constructed in this higher (or infinite) dimensional space $\mathbf{\Gamma}$. The optimization constraints are:

$$\begin{aligned} \mathbf{w}^{'}\Phi(\mathbf{x}_i) + b &\geq +1 - \xi_i, \text{if } y_i = +1 && (9.12)\\ \mathbf{w}^{'}\Phi(\mathbf{x}_i) + b &\leq -1 + \xi_i, \text{if } y_i = -1 && (9.13)\\ \xi_i &\geq 0, \forall i. && (9.14) \end{aligned}$$

In the original space of the data, $\mathbf{R}^d$, the decision boundary can be expressed as:

$$f(\mathbf{x}) = \sum_{i=1}^{N_S} \alpha_i y_i \Phi^{'}(\mathbf{s}_i)\Phi(\mathbf{x}) + b = \sum_{i=1}^{N_S} \alpha_i y_i K(\mathbf{s}_i, \mathbf{x}) + b, \tag{9.15}$$

where $K(\mathbf{x}_i, \mathbf{x}_j) = \Phi^{'}(\mathbf{x}_i)\Phi(\mathbf{x}_j)$ is a kernel function [32], [33]. It is introduced here, so that in the training algorithm Φ does not need to be explicitly known. In our experiments the Gaussian radial basis function kernel was applied, because of its reported advantages over the other types of kernel functions [58], [59], [60]. The Gaussian radial basis function kernel is expressed as:

$$K(\mathbf{x}, \mathbf{y}) = e^{-\|\mathbf{x}-\mathbf{y}\|^2/2\sigma^2} \tag{9.16}$$

Support vector machine applied to multiuser problem

Note that a support vector machine is a binary classifier. Since multiple users are involved in grip-pattern verification, we implemented a support vector machine using the method proposed in [61], which converts a multi-class problem into a binary one.

Specifically, a "difference space" was formulated, which explicitly captured the dissimilarities between two grip-pattern images. In this space we were interested in the following two classes. First, the within-class difference set, which contained the dissimilarities between grip-pattern images of the same individual. Second, the between-class difference set, which contained the dissimilarities between images of different subjects. These two classes were the input to our support vector machine algorithm, which generated a decision boundary.

9.3 Experiments, results and discussion

Prior to classification using the support vector machine, the grip patterns were transformed first by PCA (Principal Component Analysis) and then by LDA (Linear Discriminant Analysis), just as in the case of likelihood-ratio classifier. This is because we found that the verification performance would be degraded significantly, without these two transformation steps. The verification performance was evaluated by the overall equal-error rate of all the subjects. And the experiment set-up is the same as described in Section 3.3.

Note that in earlier chapters, three techniques were applied to improve the verification performance, prior to classification using the likelihood-ratio classifier. They either reduced the variations between the training and test grip-pattern images of a subject, or helped to model the variations during training of the classifier. First, we used template-matching registration (TMR) to reduce the variation between the training and test images of a subject caused by hand shift. With TMR the equal-error rate for verification was reduced from about 15% to 13% on average. The second technique that we applied was the double-trained model (DTM). Specifically, the grip patterns recorded in two out of three collection sessions were combined for training, and those of the remaining session were used for

testing. With DTM, the grip-pattern variations between the training and test images of a subject were much better modelled during training of the classifier. The verification performance was greatly improved with DTM, and the equal-error rate was reduced from about 15% to 8% on average.

Third, we applied an image processing approach prior to classification, i.e., local absolute binary patterns (LABP). With respect to a certain pixel in a grip-pattern image, the LABP processing quantifies how its neighboring pixels fluctuate. This technique reduced the pressure distribution between the training and test images of a subject. At the same time the information of hand shape was extracted from a grip-pattern image, which remained constant for a subject. The application of LABP improved the verification performance significantly, with the equal-error rate reduced from about 15% to 9% on average. When these three techniques were applied together, the verification performance was improved even more than each of them was used individually, yielding an average equal-error rate of 3% approximately.

In our experiments two parameters were optimized: the error penalty parameter C in (9.9), and the Gaussian radial basis kernel parameter σ in (9.16). The performance of grip-pattern verification was compared in three cases, using the support vector machine and the likelihood-ratio classifier respectively. In the first case, none of the three techniques of TMR, DTM and LABP was applied. The experimental results are shown in Table 9.1. "LRC" and "SVM" represent the likelihood-ratio classifier and support vector machine, respectively. The 'Average' verification results were estimated based on the matching scores, obtained from all cases of the combinations of training and test sessions. Only one of the three techniques was applied in the second case. These experimental results are shown in Table 9.2, 9.3, and 9.4. Finally, in the third case all of the three methods were applied. Table 9.5 shows the experimental results.

One can see from Table 9.1 that when none of the methods of TMR, DTM and LABP was applied, the verification performance using the support vector machine was much better on average, compared to that using the likelihood-ratio classifier. This suggests that the support vector machine was more capable of capturing the variations between the training and test grip-pattern images, than the likelihood-ratio classifier. This can be attributed to the different working principles of these two classifiers. Since the support vector machine tries to maximize the margin between different

Table 9.1: Verification results in equal-error rate (%) without TMR, DTM or LABP applied.

Train session	Test session	LRC	SVM
1	3	24.1	17.6
2	3	19.0	14.7
1	2	7.9	4.4
3	2	20.2	11.6
2	1	5.5	4.0
3	1	14.7	11.5
Average		15.0	10.2

Table 9.2: Verification results in equal-error rate (%) with TMR applied.

Train session	Test session	LRC	SVM
1	3	18.4	10.6
2	3	18.9	13.2
1	2	6.0	5.8
3	2	17.8	15.5
2	1	3.9	4.1
3	1	12.9	10.7
Average		12.7	9.6

Table 9.3: Verification results in equal-error rate (%) with DTM applied.

Train session	Test session	LRC	SVM
1+2	3	13.7	12.7
1+3	2	5.7	5.3
2+3	1	4.0	4.5
Average		7.8	7.2

classes, it has better generalization performance than the likelihood-ratio classifier, which is based on estimation of data's probability density [34], [35], [36], [37], [38].

Table 9.2, 9.3, and 9.4 show that the verification performance of both classifiers became improved on average, when TMR, DTM, or LABP was

Table 9.4: Verification results in equal-error rate (%) with LABP applied.

Train session	Test session	LRC	SVM
1	3	10.0	6.9
2	3	16.7	9.8
1	2	5.2	3.1
3	2	11.9	10.3
2	1	4.8	3.6
3	1	8.6	8.8
Average		9.2	7.2

Table 9.5: Verification results in equal-error rate (%) with TMR, DTM and LABP all applied.

Train session	Test session	LRC	SVM
1+2	3	4.9	7.6
1+3	2	3.6	5.8
2+3	1	2.0	3.7
Average		3.2	5.8

applied. And, the verification results based on the support vector machine were still better than those based on the likelihood-ratio classifier. However, one can also see that the verification performance of the support vector machine was now not as much superior to that of the likelihood-ratio classifier, as in the first case. This suggests that the likelihood-ratio classifier benefited more from the techniques (TMR, DTM and LABP), than the support vector machine did. It is very interesting to see that the likelihood-ratio classifier actually even outperformed the support vector machine, when TMR, DTM and LABP were applied together, as shown in Table 9.5.

To summarize, the support vector machine seemed to have lost its superiority over the likelihood-ratio classifier, once the variations between training and test grip-pattern images were reduced or better modelled during training of the classifier. And, the more these variations were reduced, or the better they were modelled during training, the more superiority the support vector machine tended to lose.

9.4 Conclusions

The performance of grip-pattern verification has been compared using the support vector machine and the likelihood-ratio classifier, respectively. It has been shown that the verification performance of support vector machine was much better than that of the likelihood-ratio classifier, given large variations between the training and test grip-pattern images for a subject. That is, it seems that the support vector machine captured the characteristics of grip-pattern data much better, compared to the likelihood-ratio classifier. However, once the variations have been reduced or modelled during training of the classifier, the support vector machine tended to lose its superiority and the performance of the likelihood-ratio classifier became even better.

Chapter 10

Grip-pattern verification with data for training and enrolment from different subjects

Abstract. *In this chapter the verification performance of grip-pattern recognition is investigated in a realistic situation, where the data for training and for enrolment come from different groups of subjects. Experimental results show that the false-acceptance rate at the target false-rejection rate equal to* 10^{-4}*, obtained in the realistic situation, is much higher than that in the case where the data for training and enrolment come from the same group of subjects. This is due to the mismatch between the properties of grip patterns from different groups of subjects. Also, it has been shown that the grip patterns of police officers are more stable, than those of the untrained subjects. This benefit, however, will disappear using the techniques that reduce the data variations across sessions.*

10.1 Introduction

All the experiments for grip-pattern verification described in previous chapters were done such, that the grip patterns used for training and those used for enrolment came from the same group of subjects. In practice, how-

ever, these two sets of data should be from different subjects. There were two reasons that we did not separate them previously. First, in the earlier stage of research we mainly focused on developing the verification algorithms. Second, the grip patterns recorded from the police officers alone were not enough for training the classifier, due to the limited number of subjects (see Chapter 2).

In order to investigate the verification performance of grip-pattern recognition in a realistic situation, where the data for training and for enrolment came from different groups of subjects, we did new experiments. Unlike those described in previous chapters, which were based on the grip patterns recorded from the police officers alone, in the new experiments we used *all* the data, collected from both the police officers and untrained subjects. The reader is referred to Chapter 2 for the description of data collection. In this way, more grip patterns were available for training, so that their characteristics could be better modelled. Also, it provided the possibility to answer one of our research questions, i.e. whether the grip patterns of the police officers are indeed more stable than those of the untrained subjects. In the experiments the likelihood-ratio classifier was used for grip-pattern recognition, as described in Section 3.2.

This chapter presents the experimental results of grip-pattern verification in the realistic situation, where the data for training and for enrolment came from different groups of subjects. The remainder of this chapter is organized as follows. Section 10.2 presents and discusses the experimental results. Afterwards, conclusions will be given in Section 10.3.

10.2 Experiments, results and discussion

Experiment set-up

We did two experiments of grip-pattern verification, with the data for training and for enrolment obtained from different groups of subjects. The first experiment explored the verification performance of grip-pattern recognition in the realistic situation for all users; whereas the second experiment investigated it with only the police officers considered as genuine users.

First experiment

In the first experiment, the verification performance was evaluated based on all the matching scores obtained from 120 runs. In each single run, six subjects were randomly selected and their grip patterns were used for enrolment and testing. The data of the remaining subjects were used for training. In this way, the data for training and for enrolment came from completely different groups of subjects. Also, enough test images from the genuine users were made available to estimate the false-acceptance rate at the target false-rejection rate of the system, i.e. 10^{-4}.

As described in Chapter 2, the grip patterns of both the police officers and untrained subjects were recorded in three sessions. In each run of the experiment, of the data from the six subjects selected for enrolment and testing, we used those of two out of three collection sessions for enrolment and the grip patterns of the remaining session for testing. This was a variant form of the double-trained model (see Section 3.3), and we named it "double-enrolled model". Also, the other methods that improved the verification performance of grip-pattern recognition were applied. They were template-matching registration, local absolute binary patterns, and maximum-pairwise comparison. The reader is referred to Chapter 4, Chapter 5 and Chapter 6 respectively for descriptions of these methods. The template-matching registration was capable of reducing the variation between the training and test images of a subject, caused by the hand shift. The local absolute binary patterns processing proved able to extract the information of hand shape from a grip-pattern image, which remained constant for a subject. Also, it can reduce the variation of pressure distribution, between the training and test images of a subject. And, a much lower false-acceptance rate at the target false-rejection rate equal to 10^{-4} can be obtained using maximum-pairwise comparison, than mean-template comparison.

Earlier in Chapter 4 a fused classifier was described, which was based on both the grip patterns preprocessed by template-matching registration, and the hand shifts produced by maximum-matching-score registration. This classifier proved to have better verification performance, than the one based on grip patterns alone. In the new experiments, however, it has not been implemented. The reason was twofold. First, running the experi-

Table 10.1: Verification results with the grip patterns for training and enrolment from the same group or different groups of subjects.

	Equal-error rate (%)	FAR*ref* (%)
RF	3.6	45
RS, all users	5.0	68
RS, police	6.1	62

ment proved to be rather time consuming, since the verification decision was based on the outputs of two classifiers. Second, we found that the verification performance did not improve much by using the fused classifier, when template-matching registration, local absolute binary patterns, and maximum-pairwise comparison were applied.

Second experiment

The second experiment was done in almost the same way as the first one, except for two changes. First, in each single run the six subjects for enrolment and testing were not *completely* randomly selected. Specifically, three of them were the police officers, and the other three were untrained subjects. Second, the verification performance was evaluated based on only the matching scores, produced by the comparison of test images and the template images from *the police officers*. The verification performance was evaluated based on all the matching scores obtained from 240 runs, so that enough test images from the genuine users were available to measure the false-acceptance rate at the false-rejection rate equal to 10^{-4}.

Experimental results and discussion

The experimental results are shown in Table 10.1. Both the equal-error rate and the false-acceptance rate at the false-rejection rate equal to 10^{-4}, denoted by 'FAR*ref*', are presented. 'RS' represents the realistic situation. As a reference the experimental results, with the grip patterns used for training and for enrolment coming from the same group of subjects, are also listed, denoted as 'RF'. The false-acceptance rate and false-rejection rate curves are shown in Figure 10.1.

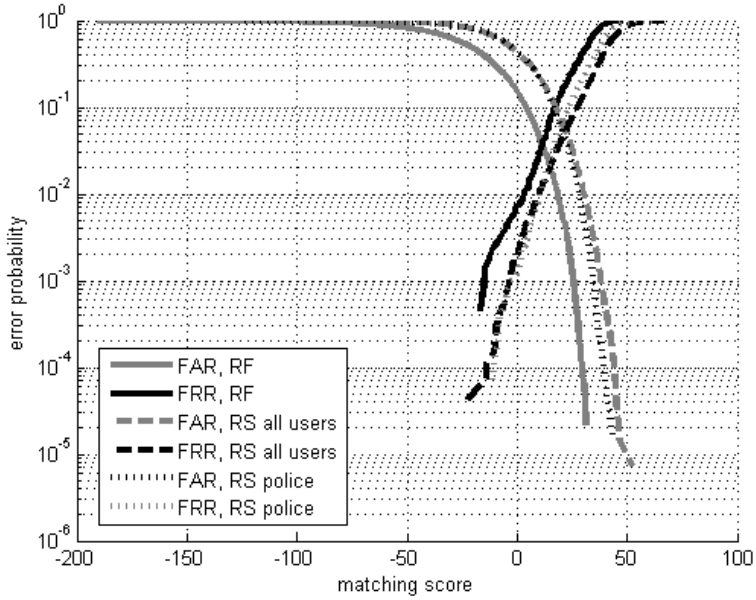

Figure 10.1: False-acceptance and false-rejection rate curves in reference case and realistic situation.

One can see that compared to the reference case, the false-acceptance rate at the false-rejection rate equal to 10^{-4} increased significantly in the realistic situation, where the grip patterns for training and for enrolment came from different groups of subjects. Figure 10.1 shows that the false-acceptance rate increased in the realistic situation, compared to that in the reference case, whereas the false-rejection rate decreased. This may be explained as follows. Since the grip patterns for training and for enrolment came from two separate groups of subjects, it was likely that the global distributions of these two sets of data were different. Therefore, the directions with the greatest variances of the total data, determined by the transformation of principal component analysis, may also be quite different. That is, in those dimensions where different classes of the grip patterns for training were distinguishable, the different classes of the data for enrolment might not. Thus, a test image from an imposter might be easily

Table 10.2: Verification results based on grip patterns from the police officers and untrained subjects.

	Equal-error rate (%)		FAR*ref* (%)	
	Police	Untrained	Police	Untrained
Within	0.6	1.2	6	13
Across, without techniques	15.0	34.1	80	95
Across, with techniques	3.6	3.4	45	48

misclassified as from the genuine user. As a result, the false-acceptance rate increased. The reason that the false-rejection rate decreased might be that some samples from the genuine users, which were misclassified as from the imposters in the reference case, were correctly classified in the realistic situation. This was attributed to the fact that after the transformation of principal component analysis, the different classes of grip patterns might become less separable, compared to the reference case. One can see that the false-acceptance rate at the false-rejection rate equal to 10^{-4} increased significantly, since the effect of the increase of the false-acceptance rate was higher than the decrease of the false-rejection rate.

One can also see that the verification performance of grip-pattern recognition was almost the same, in the case where only the police officers were considered as genuine users, as that in the general case. This seemed to indicate that the grip patterns of police officers are not necessarily more stable, than those of the untrained subjects. In order to validate this conclusion we did the within-session and across-session experiments, with the grip patterns recorded from the untrained subjects alone. In the within-session experiment none of the techniques were applied, which proved to be able to improve the verification performance by reducing the data variations across sessions, namely, double-trained model, template-matching registration and local absolute binary patterns. The across-session experiment was done in two cases, with and without these techniques applied, respectively. Table 10.2 shows the averaged experimental results. As a reference, the verification results based on the grip patterns recorded from the police officers alone are also presented.

One can see from Table 10.2 that the verification performance based on

the grip patterns of police officers were not consistently better, than that based on the data of the untrained subjects. Specifically, in the within-session case, the verification performance based on the data of the police officers was better. This was consistent with our observation described in Section 3.3. In the across-session case, where none of the techniques above was applied, the verification performance was also better based on the data of the police officers. However, in the across-session case where those techniques were applied, the verification performance based on the grip patterns recorded from the untrained subjects was comparable, to that based on the data from the police officers.

To summarize, the grip patterns of police officers are indeed more stable than those of the untrained subjects. This benefit, however, will disappear using the techniques that reduce the data variations across sessions.

10.3 Conclusions

The verification performance of grip-pattern recognition in a realistic situation was investigated, where the data for training and for enrolment came from different groups of subjects. It was shown that the false-acceptance rate at the target false-rejection rate equal to 10^{-4} became much higher, compared to the case where the grip patterns for training and for enrolment came from the same group of subjects. This may have been caused by the mismatch between the global distribution of grip patterns for enrolment, and that of the grip patterns for training.

Also, the verification performance of grip-pattern recognition was shown to be almost the same, in the case where only the police officers were considered as genuine users, as that in the general case. Combining this and our observation from other experiments, we found that the grip patterns of police officers are indeed more stable than those of the untrained subjects. However, this benefit will disappear using the techniques that reduce the data variations across sessions.

Chapter 11

Summary and conclusions

11.1 Summary

The operation of guns by others than the rightful users may pose a severe safety problem. In particular, casualties occur among the police officers, whose guns are taken during a struggle and used against themselves. A solution to this problem is the application of a smart gun that can authorize its rightful user. An important criterion is how personalized the verification system of this gun is.

In our work the verification performance of a biometric recognition system based on grip patterns, as part of a smart gun for use by the police officers, has been investigated. The biometric features are extracted from a two-dimensional pattern of the pressure, exerted on the grip of a gun by the hand of a person holding it. Such a grip-pattern verification system is interesting, particularly because it is well personalized. For this application a very low false-rejection rate is required, rendering it highly unlikely that the rightful user is not able to fire the gun. We have set the false-rejection rate for verification as below 10^{-4}, the official requirement of the failure probability of a police gun in the Netherlands. Under this precondition, the false-acceptance rate has to be minimized.

We would like to answer the following research questions. First, whether and to what extent the grip patterns are useful for identity verification.

Second, whether the grip patterns of police officers are, as expected, more stable than those of the untrained subjects. Third, whether grip-pattern recognition can be used for a police gun. Here we set the acceptable false-acceptance rate, at the false-rejection rate at 10^{-4}, to be lower than 20%.

We investigated biometric verification on the grip-pattern data collected from a group of police officers in three sessions with a time lapse, using a likelihood-ratio classifier. The likelihood-ratio classifier was selected because it is optimal in the Neyman-Pearson sense and, therefore, well-suited for our requirement. In the within-session case, where the grip patterns for training and testing were recorded in the same session, the verification performance was fairly good. In the across-session case, however, the verification performance was degraded strongly, due to the variations between the training and test images of the same subject collected in different sessions. Basically, two types of variations were observed. First, a variation of pressure distributions occurred between the grip patterns of a subject across sessions. Second, another type of variation resulted from the hand shift of a subject. However, despite the variations of grip patterns, we also found that the hand shape, as it occurred in ones grip-pattern image, was almost constant across sessions. This is easy to understand, as the hand shape of a subject is a physical characteristic that does not change so rapidly.

Based on the characteristics of grip-pattern images, the verification performance has been improved in the following three ways: modelling the data variations during training of the classifier, reducing the data variations across sessions, and extracting information of the hand shapes from images. First, we combined the grip-pattern data of two out of three collection sessions for training, and used those of the remaining session for testing. In this way, the across-session data variations were much better modelled during training of the classifier.

Second, in order to reduce the data variation caused by the hand shift, we applied template-matching registration (TMR) as preprocessing prior to classification. For comparison, the maximum-matching-score registration (MMSR) was also implemented. In MMSR a measured image is aligned such, that the matching score between this image and the template for classification attains its maximum. It was found that TMR was able to effectively improve the verification performance; while MMSR was not. How-

ever, the hand shift measured by MMSR was proved particularly useful in discriminating imposters from the genuine users. This has inspired the application of a fused classifier, based on both the grip pattern and hand shift obtained after registration.

Third, a novel approach, local absolute binary patterns (LABP), as image preprocessing prior to classification has been applied. With respect to a certain pixel in an image, the LABP processing quantifies how its neighboring pixels fluctuate. It has been proved that not only was this technique able to reduce the across-session variation of the pressure distribution in the images, but also, it was capable of extracting information of the hand shape from an image.

In this work, we have also compared two types of comparison methods for grip-pattern verification. One was mean-template comparison, where the matching score between a test image and a subject was computed by comparing this test image to the mean value of the subject's training samples. The other one was maximum-pairwise comparison, where the matching score between a test image and a subject was selected as the maximum of all the similarity scores, resulting from the comparisons between this test image and each training sample of the subject. It has been found that the major advantage of using maximum-pairwise comparison over mean-template comparison is, that a much lower false-acceptance rate at the required false-rejection rate of the system, i.e. 10^{-4}, can be obtained.

In data collection sessions, we observed that occasionally horizontal or vertical lines happened to be missing in a grip-pattern image. That is, the pressure values in those lines were all recorded as zeros. This was found to have been caused by some cable damage in the prototype, used to record the data. Since in practice there can be various factors causing missing lines in a grip-pattern image while using a smart gun, such as defects of the hardware, restoration of the missing lines is necessary. We have, therefore, presented a restoration method based on null-space error minimization. It has proved to be able to effectively restore the missing lines in a grip-pattern image.

A number of alternative approaches have also been investigated for grip-pattern verification. First, we have investigated grip-pattern verification based on two types of high-level features. One was the features measured in the correlation image between a test image and a part of the grip pattern

in an image as a template. The other one was the contour of the finger patterns in an image. Experimental results have shown that the verification performance based on high-level features was generally not good. This was due to the fact that the high-level features extracted from the grip-pattern images were not discriminative enough for identity verification.

Second, verification performance of the system using the Support Vector Machine classifier has been evaluated. It has been shown that the support vector machine has better verification performance than the likelihood-ratio classifier, if there are considerable variations between the training and test images of a subject. However, once these variations are reduced or better modelled during training of the classifier, the support vector machine tends to lose its superiority.

Finally, the verification performance of grip-pattern recognition has been investigated in a realistic situation, where the data for training and for enrolment come from different groups of subjects. It has been shown that the false-acceptance rate at the false-rejection rate equal to 10^{-4}, obtained in the realistic situation, is much higher compared to that in the case where they come from the same group of subjects. This is due to the mismatch between the properties of grip patterns from different groups of subjects. Also, one of our research questions has been answered. It has been shown that the grip patterns of police officers are indeed more stable, than those of the untrained subjects. This benefit, however, will disappear using the techniques that reduce the data variations across sessions.

11.2 Conclusions

General conclusions

- In general, the grip patterns contain useful information for identity verification. The essential requirement to the verification system as part of a smart gun, however, has not been satisfied. That is, the false-acceptance rate at the false-rejection rate equal to 10^{-4} is not low enough to be accepted. We conclude, therefore, that based on the approaches and methods investigated grip-pattern recognition is not suitable to be used for a police gun.

- If the grip-pattern data for training and testing come from the same collection session, the verification performance is fairly good. However, the verification performance becomes much worse if the grip patterns for training and testing are recorded in different sessions, respectively. This is due to the data variations across sessions.

- The grip patterns of police officers are more stable than those of the untrained subjects. This benefit, however, will disappear using the techniques that reduce the data variations across sessions.

Conclusions regarding specific approaches

- The application of double-trained model, i.e. combining the grip patterns recorded in different sessions for enrolment, has greatly improved the verification performance of grip-pattern recognition. This is because the data variations across sessions have been better modelled. Since in many biometric applications large variations occur to the data recorded with a time lapse, the method of double-trained model can be used in these applications as well.

- Template-matching registration has improved the verification performance of grip-pattern recognition. The improvement is, however, not very large due to a large variation of the pressure distribution between a measured grip-pattern image, and the template to which it is compared. This observation suggests that in a more general case of image registration if, besides the shift, there are also other variations between images, the effect of template-matching registration may be largely degraded.

- A significant improvement of the verification performance of grip-pattern recognition has been achieved by LABP. This is attributed to two effects of the LABP processing on grip-pattern images. First, the across-session variation of pressure distribution is greatly reduced. The reason is that in each image after LABP the pressure values in different subareas within the hand part become much more equalized,

compared to the case in the original image. Second, with LABP processing the information of hand shapes is extracted by enhancing the contrast between the grip patterns and the background. Therefore, after LABP the grip pattern becomes much more distinctive with respect to the background. The second effect of LABP on the grip-pattern images is similar to that of the logarithm transformation. Attributed to its first effect, however, the verification performance has been improved much more by LABP processing than by logarithm transformation.

- Maximum-pairwise comparison may improve the verification performance of a biometric application, where a low false-rejection rate is particularly required. Specifically, at a certain value of the false-rejection rate a lower false-acceptance rate may be reached.

- The verification performance of grip-pattern recognition based on high-level features is generally not good. This is due to the fact that the high-level features extracted from the grip-pattern images are not discriminative enough for identity verification, compared to the low-level features. However, in other biometric applications, where the geometry of the object to be recognized is well distinguishable, verification by high-level features may be a good option.

- The support vector machine classifier captures the characteristics of grip-pattern data fairly well, given large variations between the training and test sets. However, once the variations have been reduced or better modelled during training of the classifier, the support vector machine tends to lose its superiority. This observation of the support vector machine, we believe, is also valid in many other applications of pattern recognition.

11.3 Recommendations for further research

- We have designed the pressure sensor, which is customized to the grip of the gun. However, it has not yet been implemented by the company Twente Solid State Technology (TSST). Grip-pattern images recorded with the customized sensor are expected to have better resolution in subareas, where the data play a more important role than others for verification, and thus bring improvement of the verification performance.

- The nature of grip-pattern variation of the pressure distribution should be further investigated. Based on it, more methods may be developed to reduce this type of variation of the images. Particularly, since the local absolute binary patters (LABP) proved to work well dealing with it, more research may be done to have better understanding of this method.

- In our work it has been assumed that the covariance matrices of all users are exactly the same. The common covariance matrix has been, therefore, estimated using the grip patterns of all users. This assumption, however, may not be true. In the future, research can be done with the covariance matrix of each user estimated individually, with his or her own data.

- Grip-pattern recognition may be used for other applications, where the requirement of a low false-rejection rate is not as stringent as it is for a smart gun. Consider an intelligent door handle, for example, equipped with grip-pattern recognition and capable of recognizing who enters or leaves a room.

Acknowledgements

This thesis could not have been completed without the help and support of many people. First, I would like to express my deep gratitude to my supervisor Raymond Veldhuis. He has encouraged and helped me to overcome the difficulties in my research, especially when I felt doubtful. More important, I have learned from him the way of thinking openly and creatively. This will benefit me for the rest of my life in all aspects.

Second, I am grateful to my promoter Kees Slump, for having provided me the freedom to steer my work on my own. This has allowed me to perform research in different directions.

Third, I would like to thank my partners of this project. Wessel and Andries, your dedication to data collection has enabled the feasibility of the experiments, and made the research conclusions convincing. Here I would like to say, thank you for those joyful tours, to either a data collection session in a police station or a prototype demonstration in London.

I am also grateful to my colleagues Gert, Joost and Bas. Gert, thank you for the happy time we have spent together in our cozy office. Many conversations were so interesting, that they will sure stay in my mind and make me laugh in the future. Joost, thanks for the help with my work, especially at the beginning of it. This has enabled me to get a hang of the project within a short period. Bas, thank you for all the discussions. They have always provided me with a fresh perspective to the problem, and inspired me with new ideas to solve it.

Furthermore, my gratitude goes to the other colleagues in the Signals and Systems group. Your friendliness has created a nice relaxed atmosphere for me in these four years. This has helped me to focus on my research and accomplish it successfully. Also, thank you all for the pleasant time that

we have shared together on occasions like ProRISC.

And, I would like to thank the members of my defense committee, for spending their time reviewing my thesis. Their comments have not only helped to clarify many points in it, but also improve its readability.

Uncle, aunty and Huan, being in a country where you are too, I always feel secured and peaceful.

Chun, Shuo, thank you for always being there for me, for anything.

Geert, thank you for your love, and patience, and everything.

Finally, this thesis is dedicated to my dear pa and ma. No matter how far I go, your love is always around.

Bibliography

[1] The national Uniform Crime Reporting Program. Law enforcement officers killed and assaulted. *Federal Bureau of Investigation*, 2001.

[2] Concept Development Corporation. *http://members.aol.com/saftblok/*, 2008.

[3] J. W. Wirsbinski. Smart gun technology update. Technical report, Sandia National Laboratories, 2001.

[4] Smart Lock Technology Inc. *http://www.smartlock.com/*, 2008.

[5] Metal Storm Limited. Annual report. 2000.

[6] D. R. Weiss. Smart gun technology project final report. Technical report, Sandia National Laboratories, 1996.

[7] A. K. Jain, L. Hong, and S. Pankanti. Biometric identification. *Communications of the ACM*, 43(2):90–98, 2000.

[8] D. Maltoni, D. Maio, A. K. Jain, and S. Prabhakar. *Handbook of Fingerprint Recognition*. Springer, 2003.

[9] H. B. Adams. Safety trigger. *U.S. Patent No. 5,603,179*, feb 1997.

[10] G. T. Winer. Gun security and safety system. *U.S. Patent No. 5,459,957*, oct 1995.

[11] Z. Chen and M. Recce. Handgrip recognition. *Journal of Engineering, Computing and Architecture*, 1(2), 2007.

[12] M. Golfarelli, D. Maio, and D. Maltoni. On the error-reject trade-off in biometric verification systems. *IEEE Trans. Pattern Anal. Mach. Intell.*, 19(7):786–796, 1997.

[13] A. K. Jain and N. Duta. Deformable matching of hand shapes for user verification. In *International Conference on Image Processing (ICIP (2))*, pages 857–861, 1999.

[14] R. Sánchez-Reillo, C. Sanchez-Avila, and A. González-Marcos. Biometric identification through hand geometry measurements. *IEEE Trans. Pattern Anal. Mach. Intell.*, 22(10):1168–1171, 2000.

[15] R. L. Zunkel. *Hand Geometry Based Verification*, chapter 4, pages 87–101. Biometrics - Personal Identification in a Networded Society. Kluwer Academic Publishers, 1999.

[16] R. N. J. Veldhuis, A. M. Bazen, W. Booij, and A. J. Hendrikse. Hand-geometry recognition based on contour parameters. In A. K. Jain and N. K. Ratha, editors, *SPIE Biometric Technology for Human Identification II*, pages 344–353, mar 2005.

[17] W. L. Maness. Pressure and contact sensor for measuring dental occlusion. *U.S. Patent No. 4,856,993*, aug 1989.

[18] M. Recce. Unauthorized user prevention device and method. *U.S. Patent No. 6,563,940*, may 2003.

[19] T. Chang, Z. Chen, B. Cheng, M. Cody, M. Liska III, W. Marshall, M. Recce, D. Sebastian, and D. Shishkin. Enhancing handgun safety with embedded signal processing and dynamic grip recognition. In *The 31st Annual Conference of the IEEE Industrial Electronics Society (IECON)*, pages 2107–2113, 2005.

[20] J. A. Kauffman, A. M. Bazen, S. H. Gerez, and R. N. J. Veldhuis. Grip-pattern recognition for smart guns. In *14th Annual Workshop on Circuits, Systems and Signal Processing (ProRISC)*, pages 379–384, 2003.

[21] R. N. J. Veldhuis, A. M. Bazen, J. A. Kauffman, and P. H. Hartel. Biometric verification based on grip-pattern recognition. In *Security, Steganography, and Watermarking of Multimedia Contents*, pages 634–641, 2004.

[22] X. Shang, R. N. J. Veldhuis, A. M. Bazen, and W. P. T. Ganzevoor. Algorithm design for grip-pattern verification in smart gun. In *16th Annual Workshop on Circuits, Systems and Signal Processing (ProRISC)*, pages 674–678, 2005.

[23] X. Shang, K. Kooi, and R. N. J. Veldhuis. Restoration of missing lines in grip patterns for biometrics authentication on a smart gun. In *2nd Annual IEEE BENELUX/DSP Valley Signal Processing Symposium (SPS-DARTS)*, pages 95–98, 2006.

[24] X. Shang and R. N. J. Veldhuis. Registration of hand-grip pattern in smart gun. In *17th Annual Workshop on Circuits, Systems and Signal Processing (ProRISC)*, pages 192–195, 2006.

[25] X. Shang and R. N. J. Veldhuis. Local absolute binary patterns as image preprocessing for grip-pattern recognition in smart gun. In *IEEE First International Conference on Biometrics: Theory, Applications and Systems (BTAS)*, pages 1–6, 2007.

[26] X. Shang and R. N. J. Veldhuis. Grip-pattern verification for a smart gun. *J. Electron. Imag.*, 17(1), 2008.

[27] X. Shang and R. N. J. Veldhuis. Comparison of grip-pattern recognition in smart gun based on likelihood-ratio classifier and support vector machine. In *International Conference on Image and Signal Processing (ICISP)*, pages 289–295, 2008.

[28] X. Shang and R. N. J. Veldhuis. Grip-pattern recognition in smart gun based on maximum-pairwise comparison and mean-template comparison. In *IEEE Second International Conference on Biometrics: Theory, Applications and Systems (BTAS)*, pages 1–5, 2008.

[29] H. L. Van Trees. *Detection, Estimation, and Modulation Theory.* Wiley, New York, 1968.

[30] A. M. Bazen and R. N. J. Veldhuis. Likelihood-ratio-based biometric verification. *IEEE Trans. Circuits Syst. Video Techn.*, 14(1):86–94, 2004.

[31] J. A. Kauffman. Handgrip and squeeze pattern analysis for smart guns. Master's thesis, University of Twente, 2003.

[32] V. N. Vapnik. *The Nature of Statistical Learning Theory.* Information Science and Statistics. Springer, second edition, 2000.

[33] C. J. C. Burges. A tutorial on support vector machines for pattern recognition. *Data Min. Knowl. Discov.*, 2(2):121–167, 1998.

[34] P. Bartlett and J. Shawe-Taylor. *Advances in Kernel Methods: Support Vector Learning*, pages 43–54. MIT Press, Cambridge, MA, USA, 1999.

[35] A. Buhot and M. B. Gordon. Robust learning and generalization with support vector machines. *J. Phys. A: Math. Gen.*, 34(21):4377–4388, jun 2001.

[36] M. G. Bello and G. J. Dobeck. Comparison of support vector machines and multilayer perceptron networks in building mine classification models. In R. S. Harmon, Jr. J. H. Holloway, and J. T. Broach, editors, *Proceedings of SPIE - Volume 5089, Detection and Remediation Technologies for Mines and Minelike Targets VIII*, pages 77–94, sep 2003.

[37] W. Hämäläinen and M. Vinni. Comparison of machine learning methods for intelligent tutoring systems. In *Intelligent Tutoring Systems*, pages 525–534, 2006.

[38] C.-J. Liu and H. Fujisawa. Classification and learning for character recognition: Comparison of methods and remaining problems. In *Neural Networks and Learning in Document Analysis and Recognition, First IAPR TC3 NNLDAR Workshop*, pages 1–7, Aug 2005.

[39] R. O. Duda, P. E. Hart, and D. G. Stork. *Pattern Classification.* John Wiley & Sons, second edition, 2001.

[40] W. S. Sarle. *ftp://ftp.sas.com/pub/neural/FAQ.html*, may 2002.

[41] A. K. Jain. *Fundamentals of Digital Image Processing*. Information and System Sciences. Prentice Hall, 1989.

[42] L. J. Spreeuwers, B. J. Boom, and R. N. J. Veldhuis. Better than best: Matching score based face registration. In R. N. J. Veldhuis and H. S. Cronie, editors, *Proceedings of the 28th Symposium on Information Theory in the Benelux, Enschede, The Netherlands*, pages 125–132, Eindhoven, May 2007. Werkgemeenschap voor Informatie- en Communicatietheorie.

[43] B. J. Boom, G. M. Beumer, L. J. Spreeuwers, and R. N. J. Veldhuis. Matching score based face recognition. In *17th Annual Workshop on Circuits, Veldhoven, The Netherlands*, pages 1–4, November 2006.

[44] C. Ji and S. Ma. Combinations of weak classifiers. In *Advances in Neural Information Processing Systems (NIPS)*, pages 494–500, 1996.

[45] A. A. Ross, K. Nandakumar, and A. K. Jain. *Handbook of Multibiometrics*. Springer, 2006.

[46] Q. Tao and R. N. J. Veldhuis. Optimal decision fusion for a face verification system. In *The 2nd International Conference on Biometrics (ICB)*, pages 958–967, 2007.

[47] R. N. J. Veldhuis, F. Deravi, and Q. Tao. Multibiometrics for face recognition. *Datenschutz und Datensicherheit (DuD)*, 32(3):204–214, March 2008.

[48] Q. Tao, R. T. A. van Rootseler, R. N. J. Veldhuis, S. Gehlen, and F. Weber. Optimal decision fusion and its application on 3d face recognition. In A. Bromme, C. Busch, and D. Huhnlein, editors, *Proceedings of the Special Interest Group on Biometrics and Electronic Signatures, Darmstadt, Germany*, GI-Edition, pages 15–24, Germany, July 2007. Gesellschaft fur Informatik e.V.

[49] Q. Tao and R. N. J. Veldhuis. Optimal decision fusion for verification of face sequences. In R. N. J. Veldhuis and H. S. Cronie, editors, *Proceedings of the 28th Symposium on Information Theory in the*

Benelux, Enschede, The Netherlands, pages 297–303, Eindhoven, May 2007. Werkgemeenschap voor Informatie- en Communicatietechniek.

[50] G. Heusch, Y. Rodriguez, and S. Marcel. Local binary patterns as an image preprocessing for face authentication. In *Automatic Face and Gesture Recognition (FG)*, pages 9–14. IEEE Computer Society, 2006.

[51] T. Ojala, M. Pietikäinen, and D. Harwood. A comparative study of texture measures with classification based on featured distributions. *Pattern Recognition*, 29(1):51–59, 1996.

[52] J. C. Russ. *Image Processing Handbook*. CRC Press, Taylor & Francis Group, fifth edition, 2007.

[53] R. N. J. Veldhuis. *Restoration of Lost Samples in Digital Signals*. Prentice Hall, New York, 1990.

[54] C. Cortes and V. Vapnik. Support-vector networks. *Machine Learning*, 20(3):273–297, 1995.

[55] A. Shmilovici. Support vector machines. In *The Data Mining and Knowledge Discovery Handbook*, pages 257–276. 2005.

[56] T. Evgeniou and M. Pontil. Support vector machines: Theory and applications. In *Machine Learning and Its Applications*, pages 249–257, 2001.

[57] P. W. Wagascha, B. Manderick, and M. Muuro. On support vector and relevance vector machines. In *Belgium-Netherlands Conference on Artificial Intelligence (BNAIC)*, pages 297–304, 2005.

[58] C.-W. Hsu, C.-C. Chang, and C.-J. Lin. A practical guide to support vector classification. Technical report, 2007.

[59] S. S. Keerthi and C.-J. Lin. Asymptotic behaviors of support vector machines with gaussian kernel. *Neural Computation*, 15(7):1667–1689, 2003.

[60] H.-T. Lin and C.-J. Lin. A study on sigmoid kernels for svm and the training of non-psd kernels by smo-type methods. Technical report, 2003.

[61] P. J. Phillips. Support vector machines applied to face recognition. In *Advances in Neural Information Processing Systems (NIPS)*, pages 803–809, 1998.