

Published in IET Biometrics
 Received on 7th August 2013
 Revised on 13th December 2013
 Accepted on 7th January 2014
 doi: 10.1049/iet-bmt.2013.0054

Special Issue: Integration of Biometrics and Forensics



Measuring and mitigating targeted biometric impersonation

John D. Bustard¹, John N. Carter¹, Mark S. Nixon¹, Abdenour Hadid²

¹Communications, Signal Processing and Control, School of Electronics and Computer Science, University of Southampton, Southampton, UK

²Center for Machine Vision Research, Department of Computer Science and Engineering, University of Oulu, Oulu, Finland
 E-mail: jdb@ecs.soton.ac.uk

Abstract: This study is concerned with the reliability of biometric verification systems when used in forensic applications. In particular, when such systems are subjected to targeted impersonation attacks. The authors expand on the existing work in targeted impersonation, focusing on how best to measure the reliability of verification systems in forensic contexts. It identifies two scenarios in which targeted impersonation effects may occur: (i) the forensic investigation of criminal activity involving identity theft; and (ii) implicit targeting as a result of the forensic investigation process. Also, the first partial countermeasure to such attacks is presented. The countermeasure uses client-specific Z -score normalisation to provide a more consistent false acceptance rate across all enrolled subjects. This reduces the effectiveness of targeted impersonation without impairing the systems accuracy under random zero-effort attacks.

1 Introduction

Biometric verification systems validate a user's claimed identity by comparing a stored biometric signature against their presented biometric features. Identification can be based on any physical features, with common examples including the face, fingerprints and the iris [1]. Such systems are commonly used to control access to secure areas of a building or complex but can also be used in criminal investigations to compare suspects to biometric signatures obtained at a crime scene. Currently, DNA and fingerprints are the most common sources of biometric evidence [2], but other forms of signature are also possible, such as facial images from CCTV recordings.

An automated biometric system is a valuable tool for criminal investigations and there are now commercial systems in regular use [3]. One difficulty, however, is the reliability of such systems. No verification process is perfect, and so may fail to discover a match (false negative) or report a match incorrectly (false positive). This paper is concerned with two situations in forensics that significantly increase the chances of false positives being reported. One is targeted biometric impersonation [4] and the other is as a result of individuals being preselected for assessment. Targeted biometric impersonation involves locating an innocent person in the verification system with a similar biometric signature and then fraudulently assuming that identity to spoof a verification check. This would be an attack by determined and sophisticated criminals. Unfortunately, in routine forensic investigations, the same situation is created inadvertently when someone is selected for a verification check based on initial identification from,

say, CCTV images or a photo-fit description. The reliability of the verification produced by the system is closer to that of a targeted attack, as the suspect has not been selected randomly from the population.

It should be noted that, although there has been significant research towards providing a statistical basis for validating forensic evidence, it is generally accepted that forensic examiners cannot provide certainty of verification. Their role is instead to provide expert comparison of forensic evidence. However, in order for automated biometric systems to provide useful information to a forensic examiner, their performance must be communicated in an understandable way. As such, it is important to ensure biometric verification systems are assessed in the likely context of their use, including the potential reduction in accuracy caused by targeted attacks.

Targeted impersonation was first introduced as a method of spoofing gait verification systems [5]. A recent conference paper has shown that this attack is also effective against face verification [4]. This paper is an expansion of that publication. It examines in more detail the context under which such attacks can occur and introduces the first countermeasure to reduce the effectiveness of targeted attacks.

The paper starts by surveying the existing literature on the measurement of biometric vulnerabilities. It then expands upon the two contexts in which targeted impersonation can be an issue for forensic investigations: deliberate targeted biometric impersonation and implicit targeting within a criminal investigation. The paper also discusses the similarities that exist between biometric menagerie effects and targeted impersonation. It then examines the effect of

targeted spoofing on an example face verification system. The investigation uses a publicly available biometric algorithm and data set. The paper then analyses how the effectiveness of such attacks increases with the number of potential targets. It concludes by describing the countermeasure and an additional metric for assessing verification performance.

2 Biometric vulnerabilities

Technology evaluations of biometric systems primarily measure verification performance using the false rejection and false acceptance rates of the system under test with different trade-off priorities [6].

Many contextual factors, such as facial pose and lighting, can have a significant effect on verification performance and, as the various biometrics have matured, these factors have been investigated [7]. More recently, deliberate attempts to attack biometric systems have been studied. Uludag and Jain [8] have identified eight different types of attack based on the part of the biometric system being subverted. The first of these types focuses on attacks aimed at the sensor. These attacks are the focus of this paper. The remaining types are attacks on the electronic systems and enrolment procedures used to set up and perform verification.

In terms of sensor-level attacks, three existing methods have been identified [9]:

- *Zero-effort attacks*, in which a person claims a random identity and attempts to be incorrectly accepted by the system. Zero-effort attacks are the attack type being measured in existing large-scale performance evaluations that calculate false accept rates.
- *Brute force attacks*, which repeatedly attempt to access a system, adjusting a biometric feature until a sufficiently close match is obtained [10]. Such attacks generally require unrestricted access to the biometric system (e.g. picking a biometric lock on a stolen laptop). Secure access control scenarios, such as passport control at an airport, make such attacks less feasible as access failures can raise alarms.
- *Artefact attacks*, which use a synthetic biometric feature that has been produced from a genuine user. Such attacks would also cover the attempted use of a surgically removed biometric features and methods which exploit residual features on a sensor [11].

An additional consideration is that not all the users of a system will necessarily have the same level of security. This was highlighted by Doddington and co-workers [12], who measured the relative recognisability of different users of a speaker recognition system. Here, the users were classified into four different types: *sheep* who have normal performance, *goats* who are difficult to recognise, *lambs* who are easy to impersonate and *wolves* who can easily impersonate others. Attackers can exploit this variation to compromise a biometric system. For example, a lamb insertion attack [9] would involve deliberately enrolling a person or synthetic feature that is known to have a similar signature to many subjects. The system containing the lamb subject would then be vulnerable to imposters claiming the lamb identity.

By deliberately selecting a legitimate user with similar biometric features, a targeted attack can enable imposters to greatly increase their chances of false acceptance. Targeted attacks are a significant vulnerability as they have no artificial traits that can be recognised, either by an

automated system or a human supervisor. They are also possible without control over the enrolment procedure or the need for a confederate whose true identity would be made known, as is the case for twin impersonation or lamb injection attacks. Such attacks are also quite likely, as they are a plausible strategy for even relatively unsophisticated attackers.

2.1 Deliberate targeted biometric impersonation

In January 2010, Al-Mabhouh, co-founder of the military wing of Hamas, was assassinated in Dubai. According to Dubai's authorities, there were up to 29 suspects, 12 of whom carried British passports, six Irish, four French, one German, four Australian and two Palestinian. Interpol and the Dubai police believe that the suspects stole the identities of real people [13]. This example highlights the risk that sophisticated attackers can undermine existing identification systems by targeting individuals for impersonation.

Increase in the use of social networking, online dating and centralised biometric databases have made identity systems more vulnerable to targeted attacks. These large searchable collections of face and other biometric data increase the chance of finding a target that has a closely matching biometric signature. Such attacks are particularly dangerous as they can be effective both against automated biometrics and manual methods of identification, such as visual passport inspection.

2.2 Implicit targeting within forensic investigations

As a result of the use of forensic investigation within popular entertainment, jurors now expect greater forensic evidence as the basis for criminal convictions [14]. However, such evidence can be problematic if the chances of a false match are high. In practice, because the prior probabilities are not known, forensic experts do not make direct false match claims. However, in order to provide meaningful judgement, even informally, the reliability of forensic approaches needs to be understood. Automated biometric systems offer the potential for a formal and repeatable approach to forensic comparison. However, the false acceptance rates of such systems are typically analysed within controlled academic studies that fail to capture the context of a criminal investigation. In particular, biometric verification evaluations implicitly assume that comparisons between different subjects would result from a random pairing of individuals from a population. In practice, if the basis for the comparison is transparent to investigators, the subject pairs would have been preselected because they appear visually similar. For example, such preselection would occur in the case where a suspect's face is being compared to CCTV footage of a person committing a crime. In this case, the reliability of such a match is affected by both the chance of false acceptance due to the investigators judgment and the chance of false acceptance due to the automated system verifying a subject that had been targeted as a similar match. This means that a targeted false acceptance metric provides a more conservative and more realistic measurement of matching accuracy than traditional false acceptance measures.

2.3 Targeted impersonation and menagerie effects

Targeted impersonation can be considered as a generalisation of biometric menagerie effects. Biometric menageries classify

individuals according to how similar they are to the population as a whole and how much variation they have in their own appearance. In contrast, targeted impersonation highlights whether certain groups or pairs of individuals are similar. Such similarity can occur because of menagerie effects but can also be a result of minority groups having high levels of similarity while remaining relatively different from the population as a whole. For example, the facial verification targeted false acceptance rate of a population made up of identical twins is likely to be extremely high. The same population examined for menagerie effects or with a traditional zero-effort false acceptance rate may well show relatively low error rates as each individual may be relatively distinct from all the non-twin subjects. Similar differences could occur with minority racial groups within a population. Each member of such a group may be easily distinguishable from members of other groups but relatively hard to distinguish from their own. Under a targeted impersonation evaluation, false acceptance rates could be relatively high as imposters only select targets from their own racial demographics group. However, traditional evaluations and menagerie classifications could conceal this vulnerability because they measure similarity against the population as a whole.

3 Impact evaluation

This section evaluates the effects of targeted attacks on the CSU baseline algorithm developed by Bolme *et al.* [15] for the good, the bad and the ugly face recognition challenge [16]. The system has been trained using images from the NIST multiple biometric grand challenge data set [17]. The verification system has partial robustness to lighting variation, expression changes and occlusions. However, its performance is much lower than that has been demonstrated with state-of-the-art commercial face verification algorithms [18]. The system was evaluated using the colour FERET face database, which has been available since 1996. The frontal face subset, consisting of files labelled Fa and Fb, has been selected as it is more representative of relatively controlled face verification recordings and is consistent with the original FERET verification testing protocol [19]. The data set is made up of 1009 subjects of varying age, sex and race. The evaluation assumes the attacker has complete access to the gallery of subjects and the verification algorithms used by the system. In each case, half of the recordings of each subject are randomly selected and used as the gallery to which the attacker has access.

Target selection can be defined as follows

$$x_{\text{target}}(a) \in E, \forall y \in E, \text{Ams}(ES(y), AS(a)) < \text{Ams}(ES(x_{\text{target}}), AS(a)) \quad (1)$$

where x_{target} is the identity of the target that is chosen by attacker a , E is the set of enrolled subjects, $AS(x)$ is the set of biometric signatures that the attacker x can use for target selection, $ES(x)$ is the set of enrolled samples associated with subject x and $\text{Ams}(X, Y)$ is the average match score value produced by all combinations of biometric signatures from the sets X and Y .

Each subject in the gallery takes the role of an attacker. In each case, the gallery data is analysed to select a target that the attacker will impersonate. In all of the targeted attacks, a target was chosen based on the best match score value of

all of the possible combinations of attacker and target recordings within the gallery. The non-gallery recordings of the target are then compared against the attacker to determine imposter scores. Score values are also calculated for all the true matching pairs of users of the system. These score values are used to produce detection error tradeoff (DET) curves showing the trade-off of false accept and false reject rates for different verification thresholds. A traditional zero-effort DET curve is also produced to show the relative effect of targeted attacks. The curve is calculated by comparing each of the excluded recordings against each of the gallery recordings to produce a range of scores for both legitimate and zero-effort attacks. Examples of enrolled subjects and their attackers can be seen in Fig. 1. In most of the cases, the age, sex and race of the subject is a good match given the size of the database. However, it can be seen in the bottom right example that if face shapes are close, these factors may differ significantly. It is expected that real deployments may have more challenging input data and in turn may have more sophisticated verification systems; however, the experiments indicate that the relative effect of targeting is sufficient to warrant further investigation.

Fig. 2 shows the baseline zero-effort attack DET curve and the false acceptance rates when targeting is applied at the baseline equal error rate (EER) threshold value. The EER of the baseline is 17%. However, when a targeted attack is performed on the same system, the false acceptance rate rises to 51%, three times the original value and a significant security risk. If the threshold of the system is selected with the knowledge of targeted attacks, the EER becomes 28%, which reduces the risk but increases the false reject rate to an impractical level.

3.1 Number of targets

In the baseline experiments, the number of targets available to the attacker is necessarily restricted by the size of the data sets. The size of these data sets is consistent with the number of subjects that might access a secure office environment but is much lower than many important identity scenarios such as passport control. To analyse the effect of increasing target numbers, further experiments were performed using the face verification system. The modelling procedure is described in Algorithm 1, which is included in Appendix 1. 800 gallery subsets of increasing size were created. These subsets were used in the selection of targets for evaluation. To minimise any potential bias caused by subset selection, for a given size, all non-overlapping subsets within the first 800 subjects were combined to produce median false accept rates across the different subsets. This ensures that a subset size of 1 is virtually identical to the baseline performance. All gallery members took the role of attackers using the subset to generate the imposter scores.

Fig. 3 indicates how the false accept rate increases as the size of the target subset increases. The graph shows the false acceptance rate for a threshold that achieves the equal error rate of the baseline system under zero-effort attacks. This is a plausible threshold for systems that are unaware of the risks of targeted attacks. As the number of available targets increases, the number of possible subsets decreases, increasing the error in the measured false accept rate. Much of the curve, however, conforms reasonably well to a least squares fit of an $a \log(x) + b$ model, with $a = 5.1$ and $b = 17.8$.



Fig. 1 Examples of enrolled subjects and their attackers
Each attacker is to the right of the enrolled subject

One difficulty in using a logarithmic fit to predict false acceptance rate (FAR) is that such a curve can produce values below 0% and above 100%. Although the FAR values are limited in this way, the difference between individual biometric signatures may not be. There are many different score distributions that could produce 0% or 100% FAR values based on the relative difference between legitimate and imposter score values. As such, the logarithmic fit can be seen as expressing the functional shift in the difference

between legitimate and imposter score distributions rather than the FAR value itself. As the FAR measurements approach the bounds, excessively distant or close score values will have a diminished effect on the measured FAR.

One way to understand this effect is to treat the logarithmic prediction as the centre of a probability distribution over FAR values that can pass outside the bounds. This probability distribution reflects the likelihood of obtaining any particular FAR when the biometric system is evaluated. When determining the likelihood of 100% or 0% FAR values, the entire probability distribution outside of the bounds are combined. In practice, this means that when the targeted FAR value reaches 100%, the model predicts that there is a 50% chance of obtaining 100% FAR for any given evaluation of the system. Further research is required to determine

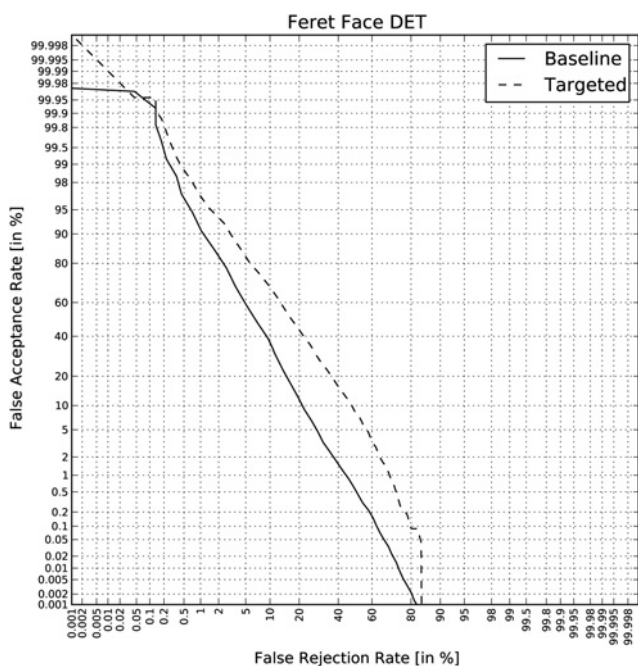


Fig. 2 Effects of a targeted attack on the CSU face verification algorithm

Baseline shows the performance of the system under a zero-effort attack. Targeted shows the increase under targeted biometric attacks

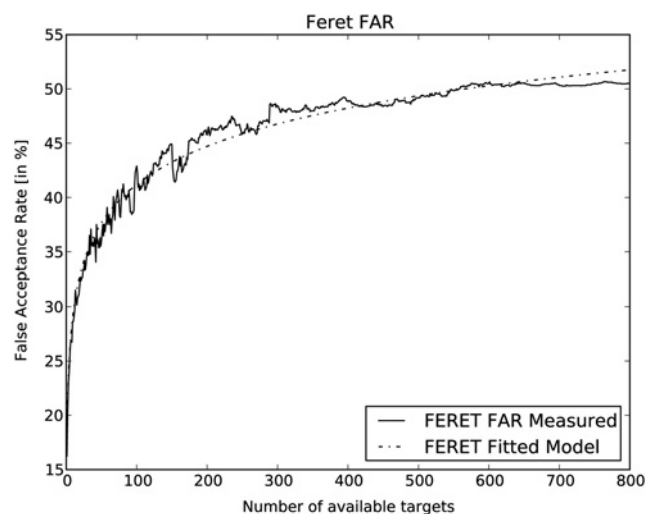


Fig. 3 Effect of target numbers on FAR with a verification threshold set at the EER of the baseline system

the shape of this distribution and to validate these predictions on systems which reach these bounds. It should be noted that this differs from typical model estimates which predict the mean FAR value. In this case, the modelled FAR value is the value of the median of the predicted measured scores and is thus not altered by the clamping effects of samples above or below this median measurement. In particular, 100% predicted FAR can occur when there are targets who have a small range of appearance variation centred on the enrolled data for each subject and thus entirely reside within the bounds needed to achieve the FRR of the system. However, it should be noted that larger evaluations are needed to confirm the accuracy of the model as it reaches 0 or 100% values.

An additional consideration is how feasible is it for attackers to obtain information about the gallery subjects and the system being attacked. For small-scale deployments, surveillance may be sufficient to establish possible targets. However, some biometrics may be more vulnerable. For example, face, voice and gait are relatively easy to record at a distance while fingerprint, iris and finger vein may require more elaborate social engineering to obtain. For identity applications with a large number of users, such as passports, public information may be sufficient. For example, a number of online dating websites have photographs of millions of users which can be anonymously searched using soft biometric constraints including, age, sex, race, hair colour and height [20]. Centralised databases of biometric information are of greater concern. For example, if the US visit database was hacked, its recordings could be used to identify possible targets for face or fingerprint attacks.

4 Targeting countermeasure

Unlike many biometric attacks, targeting has no artificial component and thus cannot be detected directly. Instead, countermeasures must attempt to limit the false acceptance rate of worst case false matches between enrolled subjects and the user population. Different enrolled subjects may have different levels of vulnerability as some subjects may be more generic than others. When targeted attacks are performed these, more generic subjects are more likely to be successfully impersonated. To minimise false acceptance rate these subjects need to have their match scores normalised so that they correspond to the likelihood that two subjects are the same.

The vulnerability of each enrolled subject can be estimated by finding the closest potential target within the enrolment database. The match score between each subject and their target provides an estimate for the false match score of an unknown user performing a targeted attack. By reducing a match score by the enrolment target score, the variation in vulnerability may be reduced. However, this assumes that the remaining database is a reasonable estimate of the population as a whole. Unfortunately, for relatively small data sets, such as the FERET face database, this estimate is insufficiently robust and provides negligible improvement in overall performance. In addition, the noise in the client-specific estimate increases the false acceptance rate of the system under zero-effort attacks.

A more reliable estimate of client-specific vulnerability can be obtained by approximating each enrolled subject's false match score distribution by a Gaussian. The mean and standard deviation of this Gaussian can then be

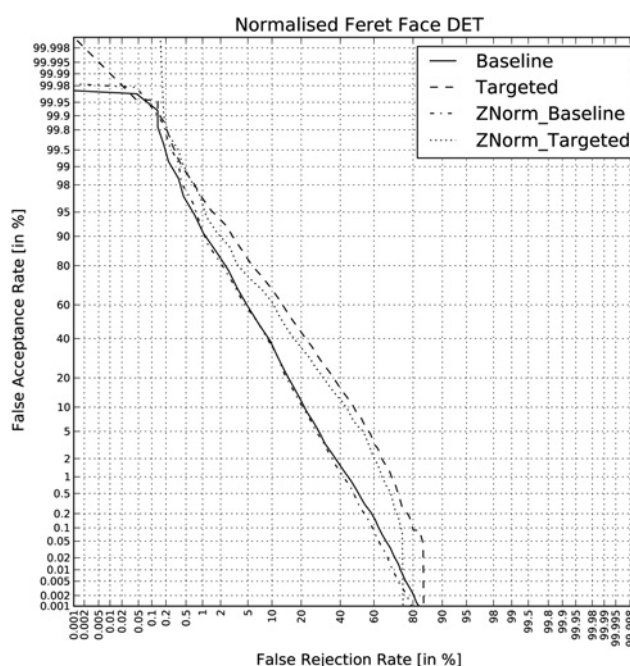


Fig. 4 Effects of a Z-score normalisation to a targeted attack on the CSU face verification algorithm

Baseline shows the performance of the system under a zero-effort attack. *Targeted* shows the increase under targeted biometric attacks. *ZNorm_Baseline* shows the performance of the system under a zero-effort attack when Z-score normalisation is applied. *ZNorm_Targeted* shows the reduction in the effect of targeted attacks when the normalisation is applied

estimated using the false match scores of each subject and all the other members of the enrolment database. Z-score normalisation can be applied for each enrolled subject to try to maintain similar false match performance for each subject. The normalisation reduces match scores by the estimated Gaussian mean and then scales scores using the reciprocal of the standard deviation, that is, $(x - \mu)/\sigma$; where x is the unnormalised score value, μ is the mean of the false match scores with all the other enrolled subjects and σ is the standard deviation of these false match scores relative to their mean. Fig. 4 shows that Z-score normalisation is effective at partially reducing the false acceptance rate for targeted attacks. Normalisation reduces the FAR at the baseline EER from 51% to 42%. The normalisation also has no significant effect on the baseline zero-effort performance. Z-score normalisation makes the assumption that the vulnerability averaged over the estimated population is proportional to the vulnerability when the subject is directly selected as a target. For larger data sets, it may be possible to improve on this performance by more accurately estimating vulnerability of those subjects likely to select the enrolled subject as a target.

5 Conclusions

This paper analyses the effect of targeted attacks that can reduce the effectiveness of biometric identity verification. It has described two contexts in which targeted attack evaluations are relevant to forensics: deliberate targeted biometric impersonation and implicit targeting within a criminal investigation. It has described the difference between targeted impersonation and biometric menagerie effects and has illustrated the problem through the

evaluation of a baseline face verification algorithm. This evaluation revealed that with 800 potential targets, attacks can increase false acceptance rates of the measured system by a factor of three, reducing security to the point that it is no longer reliable for forensic comparison. Further analysis suggests that the false acceptance rate can be estimated using a simple model that is proportional to the logarithm of the number of enrolled subjects. This model provides a means to estimate the vulnerability of systems with many users. The paper also describes the first countermeasure to targeted attacks. The countermeasure uses client-specific Z-score normalisation to provide a more consistent false acceptance rate across all enrolled subjects. When the system's threshold is set to the EER of the baseline algorithm, the countermeasure reduces the relative increase in FAR due to targeted attacks from a factor of three to a factor of two and a half. Such score normalisation already forms part of many commercial systems for the purpose of dealing with varying quality levels in enrolled data and manager effects across enrolled subjects. This paper has shown that this technique is also important in mitigating targeted attacks. In particular, as demonstrated in Fig. 4, the normalisation has negligible effect on the zero-effort attack performance and so the importance of this normalisation could be missed by those implementing such a system. As further work, it would be valuable to explore alternative score normalisation methods to assess to what extent they improve resilience to targeted attacks. It would also be useful to determine the effect of targeted attacks on state-of-the-art commercial face verification algorithms. Such systems have demonstrated 100% verification accuracy on the FERET data set [18] and many include score normalisation methods that may help to mitigate these attacks. Unfortunately, owing to their closed nature, it is not possible to analyse how such systems have been trained, and why such systems are robust or vulnerable to such attacks. To facilitate further study into these forms of vulnerability, it would be valuable for open implementations of state-of-the-art biometric systems to be made available so that more detailed and realistic performance comparisons can be made. In addition, it would be useful to analyse to what extent biometric systems are vulnerable to targeted attacks when the target is selected solely using human judgment, as this is the form of attack that would be implicitly performed within a criminal investigation.

6 Acknowledgments

This work was partially funded by the EU FP7 project TABULA RASA (257289). Figs. 2 and 4 were produced using ScoreToolKit [21].

7 References

- Jain, A.K., Flynn, P., Ross, A.A.: 'Handbook of biometrics' (Springer, Secaucus, NJ, USA, 2007)
- Nickell, J., Fischer, J.F.: 'Crime science: methods of forensic detection' (University Press of Kentucky, 1999)
- 'Morphotrust investigative solutions', <http://www.morphotrust.com/pages/1010-investigation>
- Bustard, J.D., Carter, J.N., Nixon, M.S.: 'Targeted biometric impersonation', *Int. Work. Biometrics Forensics*, 2013, **1**, pp. 1–4
- Hadid, A., Ghahramani, M., Kellokumpu, V., Pietikainen, M., Bustard, J., Nixon, M.: 'Can gait biometrics be spoofed?'. *Int. Conf. Pattern Recognition*, 2012, pp. 3280–3283
- OToole, A., Flynn, P., Bowyer, K., *et al.*: 'FRVT 2006 and ICE 2006 large-scale experimental results', *IEEE Trans. PAMI*, 2010, **32**, (5), pp. 831–846
- Jain, A.K., Li, S.Z.: 'Handbook of face recognition' (Springer, Secaucus, NJ, USA, 2005)
- Uludag, U., Jain, A.K.: 'Attacks on biometric systems: a case study in fingerprints', *Secur. Steganogr. Watermarking Multim. Contents*, 2004, **6**, pp. 622–633
- Dunstone, T., Poulton, G.: 'Vulnerability assessment', *Biometric Technol. Today*, 2011, **5**, pp. 5–7
- Martinez-Diaz, M., Fierrez, J., Galbally, J., Ortega-Garcia, J.: 'An evaluation of indirect attacks and countermeasures in fingerprint verification systems', *Pattern Recognit. Lett.*, 2011, **32**, (12), pp. 1643–1651
- Matsumoto, H., Matsumoto, T., Yamada, K., Hoshino, S.: 'Impact of artificial gummy fingers on fingerprint systems'. *SPIE Optical Security and Counterfeit Deterrence Techniques IV*, 2002, vol. 4677, pp. 275–289
- Martin, A., Przybocki, M., Doddington, G., Liggett, W., Reynolds, D.: 'Sheep, goats, lambs and wolves a statistical analysis of speaker performance in the NIST 1998 speaker recognition evaluation'. *Int. Conf. Spoken Language Processing*, 1998
- 'BBC news: Interpol puts Dubai killing suspects on wanted list', http://news.bbc.co.uk/1/hi/world/middle_east/8522595.stm
- Houck, M.M.: 'CSI the reality' (Scientific American, 2006)
- Phillips, P.J., Beveridge, J.R., Draper, B.A., *et al.*: 'An introduction to the good, the bad, & the ugly face recognition challenge problem'. 2011 IEEE Int. Conf. Automatic Face Gesture Recognition and Workshops (FG 2011), 2011, pp. 346–353
- Phillips, P., Beveridge, J., Draper, B., *et al.*: 'An introduction to the good, the bad & the ugly face recognition challenge problem'. *Int. Conf. Face Gesture*, 2011, pp. 346–353
- Beveridge, J.R., Scruggs, W.T., OToole, A.J., *et al.*: 'Overview of the multiple biometrics grand challenge'. *Int. Conf. Advances in Biometrics*, 2009, pp. 705–714
- Klum, S., Jain, A., Burge, M., Klontz, J., Klare, B.: 'Open source biometric recognition'. *Int. Conf. Biometrics: Theory, Applications and Systems*, 2013
- Moon, H., Phillips, P.: 'The FERET verification testing protocol for face recognition algorithms'. *Int. Conf. Automatic Face and Gesture Recognition*, 1998, pp. 48–53
- 'Plenty of fish dating website', <http://www.plentyoffish.com>
- Anjos, A., Marcel, S.: 'Scoretoolkit documentation'. IDIAP Technical Report, 2012

8 Appendix

See Fig. 5.

Data: t the acceptance threshold that produces the EER of a zero effort attack evaluation of the full enrollment database;

Result: The constants a and b from the $a.\log(x) + b$ TargetedFAR model

for each $x_{attacker}$ **from the set** $Enrolled_Subjects$ **do**

for $num_targets = 1$ **to the** $num_Enrolled_Subjects$ **do**

Split the remaining $Enrolled_Subjects$ into $Subsets$, a set of non-overlapping subsets with $num_targets$ elements in each subset;

for each S **of** $Subsets$ **do**

Identify the x_{target} for $x_{attacker}$ within S using equation 1;

Calculate the set MS from all of the match scores between all enrolled samples of $x_{attacker}$ and the test samples of x_{target} ;

Add the percentage of the MS values above the t to the set of FARs associated with this $num_targets$ i.e.

$FAR(num_targets)$;

end

end

end

for $num_targets = 1$ **to the** $num_Enrolled_Subjects$ **do**

Calculate $FAR_{median}(num_targets)$, the median of the $FAR(num_targets)$ subset, a single FAR value for each $num_targets$;

if $FAR_{median}(num_targets) > 0\%$ **and** $FAR_{median}(num_targets) < 100\%$ **then**

Add the point $(\log(num_targets), FAR_{median}(num_targets))$ to the set of $Regression_Points$;

end

end

Calculate the Ordinary Least Squares regression of the $Regression_Points$ values producing the gradient value a and the intersection b ;

These parameters define the $TargetedFAR = a.\log(x) + b$;

Fig. 5 Algorithm 1: pseudocode describing the procedure for calculating the TargetedFAR metric