Biometrics systems under spoofing attack:
an evaluation methodology and lessons learned

Abdenour Hadid, Nicholas Evans, Sébastien Marcel and Julian Fierrez

Abstract

Biometrics already form a significant component of current and emerging identification technologies. Biometrics systems aim to determine or verify the identity of an individual from their behavioural and/or biological characteristics. Despite significant progress, some biometric systems fail to meet the multitude of stringent security and robustness requirements to support their deployment in some practical scenarios. Among current concerns are vulnerabilities to spoofing, namely persons who masquerade as others in order to gain illegitimate accesses to protected data, services or facilities. While the study of spoofing, or rather anti-spoofing, has attracted growing interest in recent years, the problem is far from being solved and will require far greater attention in coming years. This tutorial article presents an introduction to spoofing and anti-spoofing research. It describes the vulnerabilities, presents an evaluation methodology for the assessment of spoofing and countermeasures, and outlines research priorities for the future.

I. INTRODUCTION

The provision of security can entail the protection of sensitive data, services or facilities by ensuring that only authorised persons have access. Though passwords provide some protection against illegitimate access, they are often so simple that they can be guessed or easily cracked. While offering improved security, complex passwords can be difficult to remember and consequently often 'stored' via less secure means. Furthermore, the same password is often used across multiple applications or platforms meaning a cracked password can enable a fraudster to access multiple resources.

An attractive alternative to passwords involves biometric recognition. Biometrics refer to a person's behavioural and biological characteristics such as their face, fingerprint, iris, voice, hand geometry and gait. Biometric traits can be highly discriminative yet less easily lost or stolen [1]. Despite their appeal, however, biometric systems are vulnerable to malicious attacks [2]. Among them are spoofing attacks, also referred to as presentation attacks, which refer to persons masquerading as others in order to gain illegitimate access to sensitive or protected resources. As an example, a fraudster could fool or spoof a

face recognition system using a photograph, a video or a 3D mask bearing resemblance to a legitimate individual.

Even though the threat of spoofing is now well recognised, the problem is far from being solved, thus anti-spoofing research warrants far greater attention in the future. This tutorial article introduces the problem of spoofing and related research to develop anti-spoofing solutions. The focus is an evaluation methodology for assessing both the effect of spoofing and the performance of spoofing countermeasures. A case study in face recognition is included to illustrate the application of the evaluation methodology in practice. Finally, the article also includes a summary of the lessons learned through the authors' own research and outlines a number of research priorities for the future. Most of the material is based upon anti-spoofing research performed in the scope of the European TABULA RASA research project[1] which was identified as a success story by the European Commission[2].

The presentation is self-contained and aimed at both the generally knowledgeable and non-specialist. The article aims to provide an overview of the research problem, not a comprehensive survey of the plethora of anti-spoofing techniques in the literature; such surveys can be found elsewhere, e.g. for fingerprint recognition [3], face recognition [4] and speaker recognition [5]. The intention is also to stimulate further work, particularly the development of standard metrics, protocols and datasets for evaluating progress.

The article is organized as follows. Sections II and III describe generic biometric systems and their potential vulnerabilities to spoofing. The evaluation methodology is presented in Sections IV and V first, for the assessment of vulnerabilities to spoofing and second, for the assessment of spoofing countermeasure performance. Section VI describes an example case study in face recognition. Lessons learned are discussed in Section VII together with some directions for future research before the outlook is presented in Section VIII.

II. BIOMETRICS

The term 'biometrics' is derived from the Greek words *bio* (life) and *metric* (to measure). The goal of a biometric recognition system is to determine or verify the identity of an individual from their behavioural and/or biological characteristics. Applications include criminal identification, airport checking, computer or mobile device log-in, building and critical infrastructure access control, digital multimedia rights

---

[1] http://www.tabularasa-euproject.org

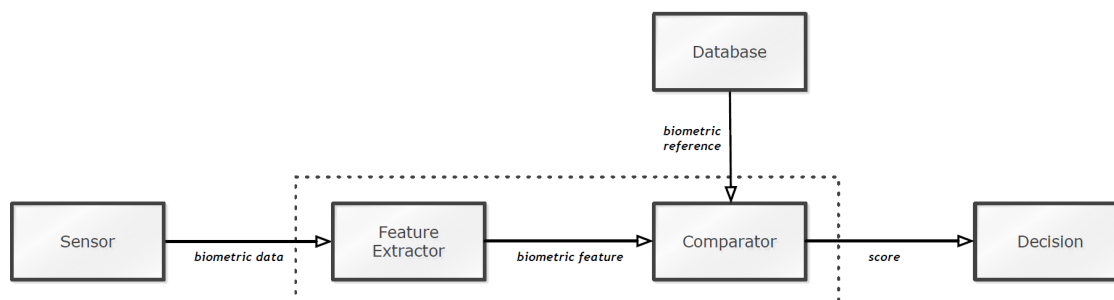[2] http://europa.eu/rapid/press-release_MEMO-13-924_en.htm

Fig. 1: A generic biometric system

control, transaction authentication, voice mail and secure teleworking. Various biometrics have been investigated, from the most conventional including fingerprint, iris, face and voice, to more emerging modalities such as gait, hand-grip, ear and electroencephalograms. Each modality has its own strengths and weaknesses [1]. For example, face recognition is among the most socially accepted biometric; face recognition is a natural method of identification used everyday by humans. In contrast, while fingerprint and iris recognition may be more reliable, they are also more intrusive. In practice the choice of biometric modality depends on the application.

Biometric systems typically function in one of two distinct modes: (i) verification (or authentication) and (ii) identification. An authentication system aims to confirm or deny a claimed identity (one-to-one matching) whereas a identification system aims to identify a specific individual (one-to-many matching). Although there are some differences between the two modes, their most basic operation, namely that of feature-to-reference comparison, is identical and consists of the following steps illustrated in Fig. 1.

First a biometric sample (e.g. a face image) is acquired from a sensor (e.g. a digital camera). Biometric features (e.g. facial intensity, colour or texture) are then extracted from the sample. These can be a set of parameters (or coefficients) which provide a compact representation of the biometric sample which is more discriminative and amenable to pattern recognition. Biometric features should minimise variations due to acquisition or environmental factors (e.g. facial expression, pose and illumination) while discriminating between the biometrics collected from different individuals.

To determine or verify the identity corresponding to a given biometric sample, the features are compared to a single (verification) or set of (identification) biometric references acquired previously during an enrolment phase. These comparisons are made by a comparator which produces a score reflecting the similarity between features and references. The decision is an acceptance or rejection in the case of verification, or the identity of the closest match in the case of identification.

(a) 3D face mask attack   (b) Fingerprint attack

Fig. 2: Example spoofing attacks: (a) face recognition spoofing using a 3D face mask; (b) fingerprint recognition spoofing using a fake fingerprint.

III. SPOOFING ATTACKS

It is now widely acknowledged that biometric systems are vulnerable to manipulation. There are two broad forms of attack: *direct* and *indirect*. Direct attacks, also referred to as *spoofing* or *presentation attacks*, are performed at the sensor level, outside the digital limits of the biometric system. Indirect attacks, however, are performed within the digital limits by intruders such as cyber-criminal hackers. These attacks may attempt to bypass the feature extractor or the comparator, to manipulate biometric references, or to exploit vulnerabilities in the communications channels.

Whereas traditional digital protection mechanisms such as encryption can be deployed to prevent indirect attacks, they cannot prevent direct attacks. Furthermore, indirect attacks require specialist expertise or equipment whereas direct attacks such as those illustrated in Fig. 2 can be implemented by the layman. Direct attacks are therefore of significant concern.

Vulnerabilities to spoofing are a barrier to the successful exploitation of biometrics technology. Unfortunately, some vulnerabilities have been exposed and well publicised in the international media. One of the earliest examples was demonstrated at Black Hat 2009, the world's premier technical security conference. Researchers from the Security and Vulnerability Research Team at the University of Hanoi (Vietnam) showed how the face recognition user authentication systems introduced by three different laptop computer manufacturers could be spoofed or bypassed using photographs of legitimate users. This vulnerability is now listed in the National Vulnerability Database maintained by the National Institute of Standards and Technology (NIST) in the US. More recently, the Chaos Computer Club, a German hacking collective, showed how an artificial finger could be used to spoof a fingerprint user authentication system developed by one of the world's most popular smart phone manufactures.
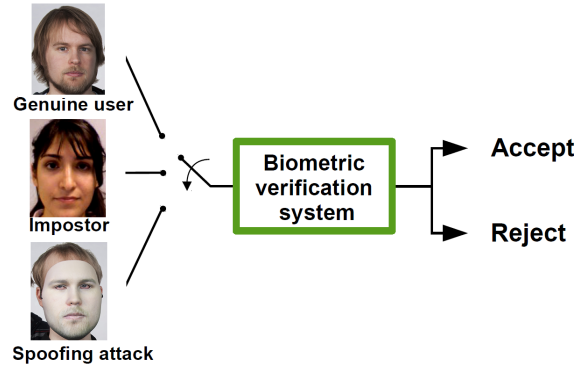
Fig. 3: A biometric verification system should accept genuine users but reject both zero-effort impostors and concerted-effort spoofing attacks.

The typical countermeasure deployed to detect spoofing attacks involves *liveness detection*, namely systems which aim to detect signs of life. Since it is inherently more difficult to spoof multiple modalities and systems simultaneously [6], multi-modal biometric systems have also been investigated as a solution to spoofing. Even so, spoofing remains a serious cause for concern, with many biometric systems remaining vulnerable even to the simplest forms of spoofing attack. Unfortunately, research to develop spoofing countermeasures is still in its infancy and warrants considerably greater attention in the future.

IV. EVALUATION METHODOLOGY: VULNERABILITIES TO SPOOFING

This section describes the first component of the proposed methodology for the evaluation of biometric systems and their vulnerabilities to spoofing. For simplicity, the following considers only biometric verification (1-to-1 matching). Traditionally, biometric systems are evaluated using large datasets of representative biometric samples with which performance is assessed according to a standard protocol of genuine and imposter trials. An evaluation essentially measures the proportion of these trials which the system misclassifies, i.e. genuine trials classified as impostor trials and vice versa. This approach, involving only casual impostors, equates to performance in the face of zero-effort spoofing attacks, i.e. where impostors make no effort to replicate the biometric traits of another, legitimate individual.

Spoofing is accomplished using fake biometric samples expressly synthesized or manipulated to provoke artificially high comparator scores. Biometric systems should thus be robust to both zero-effort impostor trials and concerted-effort spoofing attacks. This ternary scenario is illustrated in Fig. 3. The biometric system should produce high scores in the case of genuine trials, but low scores in the case of impostor and spoofed access attempts.

Assessment thus requires biometric data of three categories comprising genuine, zero-effort impostor and spoofed trials. Biometric data are typically further divided into three non-overlapping subsets: a **training set**, a **development set** and a **test set**, the purpose of which is as follows.

- The **training set** is used to train the biometric system (e.g. the learning of background models) or to develop spoofing countermeasures.
- The **development set** is used for decision threshold optimisation and the a posteriori performance at a given operating point. For each identity, the development set is partitioned into two subsets:
  - an enrolment subset used to create biometric references or models for each identity, and
  - a probe subset used for biometric comparisons (genuine, zero-effort impostor and spoofing attack trials).
- The **test set** is used to compute the a priori performance given the threshold determined from the development set. The test set is similarly partitioned into enrolment and probe subsets.

Two assessment scenarios can then be defined:

- a **licit scenario** for assessing baseline performance using genuine and zero-effort impostor trials, and
- a **spoofing scenario** for assessing vulnerabilities to spoofing.

Fig. 4 illustrates example score distributions for a licit scenario (i.e. genuine and impostor trials with no spoofing attacks). With only little overlap between the two distributions there is high potential to distinguish between genuine and impostor trials. Nonetheless, the system will still make mistakes, leading to two different error rates:

- the **False Rejection Rate (FRR)**, which reflects the proportion of genuine trials misclassified as zero-effort impostor trials, and
- the **False Acceptance Rate (FAR)**, which reflects the proportion of zero-effort impostors trials misclassified as genuine trials.

Both the FAR and FRR are dependent upon a decision threshold τ, an example of which is illustrated by the vertical, dashed line in Fig. 4. Impostor trails corresponding to a score higher than this threshold will be misclassified as genuine trials, whereas genuine trials with a score lower than this threshold will be misclassified as impostor trials.

Since the two errors are inversely related, it is often desirable to illustrate performance as a function of the threshold τ. For a specific dataset $\mathcal{D}$, one such measure is the **Half Total Error Rate (HTER)**:

$$\text{HTER}(\tau, \mathcal{D}) = \frac{\text{FAR}(\tau, \mathcal{D}) + \text{FRR}(\tau, \mathcal{D})}{2} \tag{1}$$

Fig. 4: Score distribution of genuine users (positives) and zero-effort impostors (negatives) under a licit scenario. A decision threshold τ is illustrated by the vertical, dashed line.


Fig. 5: An example detection error trade-off (DET) plot showing an equal error rate (EER) in the order of 6%.

Performance can also be illustrated with **Detection-Error Tradeoff (DET)** profiles, an example of which is illustrated in Fig. 5. DET profiles illustrate the behaviour of a biometric system for varying decision thresholds, τ, and show the trade-off between the FAR and the FRR.

In practice, the decision threshold is chosen to minimise an a posteriori performance criteria which is normally application dependent. Without focusing on a specific application, the **Equal Error Rate (EER)** is also a common reference metric. For a development dataset $\mathcal{D}_{dev}$, the EER is given by:

$$\tau^*_{\mathsf{EER}} = \arg \min_{\tau} \ |\mathsf{FAR}(\tau, \mathcal{D}_{dev}) - \mathsf{FRR}(\tau, \mathcal{D}_{dev})| \tag{2}$$

where the threshold τ^*_{EER} is set in order to equalise the FAR and FRR. The most reliable estimate of a priori performance is then determined using a test set $\mathcal{D}_{test}$ using the pre-determined threshold. For example, the HTER is determined using the decision threshold τ^*_{EER} according to:

$$\mathsf{HTER}(\tau^*_{\mathsf{EER}}, \mathcal{D}_{test}) = \frac{\mathsf{FAR}(\tau^*_{\mathsf{EER}}, \mathcal{D}_{test}) + \mathsf{FRR}(\tau^*_{\mathsf{EER}}, \mathcal{D}_{test})}{2} \tag{3}$$

Performance in the face of spoofing is assessed with the spoofing scenario, i.e. by replacing the subset of zero-effort impostor trials with spoofed trials. Fig. 6(a) illustrates example score distributions for a population of genuine, zero-effort impostor and spoofed trials. The shift between the impostor and spoofed trial distributions illustrates the likely effect of spoofing on biometric recognition performance; the overlap between the score distributions for genuine and spoofed trials is significantly greater than that between genuine and impostor distributions.

A quantitative measure of system vulnerability can be expressed in terms of the **Spoof False Acceptance Rate (SFAR)** which reflects the percentage of spoofed trials misclassified as genuine trials given a decision threshold τ. An example SFAR profile is illustrated as a function of the threshold τ in Fig. 6(b).

The vulnerability to spoofing is thus reflected in the difference between the SFAR (spoof scenario) and the FAR (licit scenario), again as a function of τ. These differences are illustrated in the example DET plots for both licit (FAR vs. FRR) and spoof (SFAR vs. FRR) scenarios in Fig. 7. Finally, it is often of interest to express system vulnerability for a specific operating point defined by a given FRR (e.g. the EER). System vulnerability is then given by the difference between the FAR and the SFAR for the same FRR (vertical, dashed line in Fig. 7). In this illustrative example for an FRR in the order of 6%, an FAR of under 10% increases to an SFAR of over 60%. This system would misclassify approximately two in three spoofed trials as genuine accesses.

V. EVALUATION METHODOLOGY: SPOOFING COUNTERMEASURES

This section describes the extension of the evaluation methodology to assess spoofing countermeasures. Fig. 8(a) illustrates the deployment of spoofing countermeasures as independent sub-systems. Just like
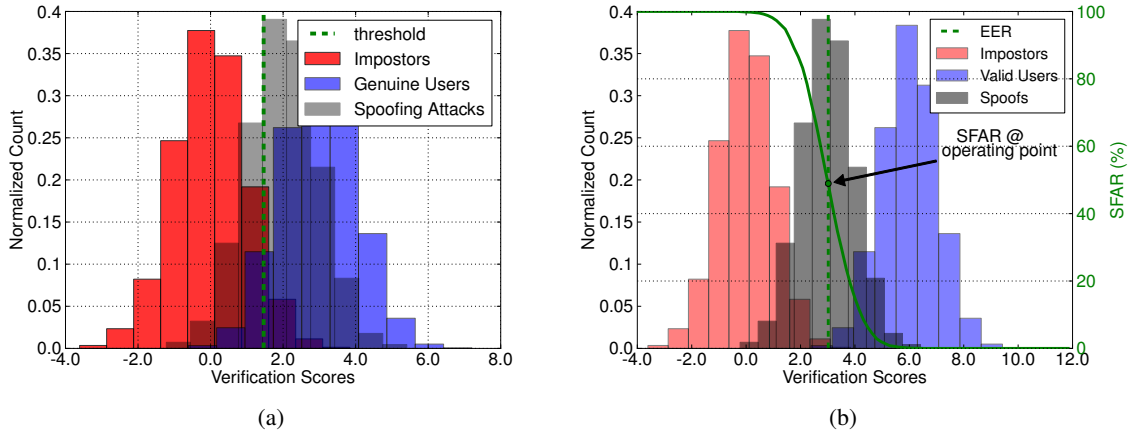
Fig. 6: Example score distributions: genuine users, zero-effort impostors, and spoofing attacks with the Spoof False Acceptance Rate (SFAR) as a function of the decision threshold τ.


Fig. 7: A detection-error tradeoff (DET) profile illustrating the Spoof False Acceptance Rate (SFAR). FRR@EER refers to the false rejection rate (FRR) on the test set for a decision threshold τ for which the FAR and FRR are equal (determined from the development set).

the biometric system, the countermeasure is a two-class classifier, its role being to distinguish between genuine and spoofed trials. Again, just like the biometric system, it will also make mistakes, leading to two new error rates referred to as the **False Living Rate (FLR)** and the **False Fake Rate (FFR)**. The

(a) Two independent components

(b) Fused sub-systems
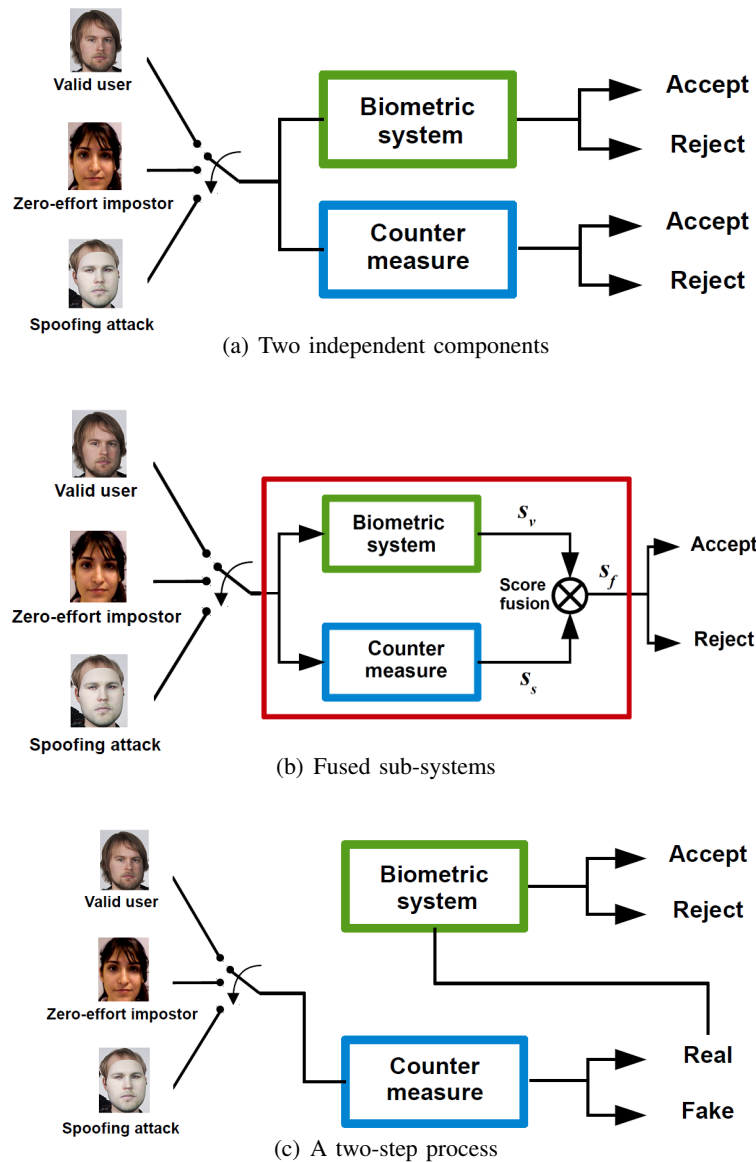

(c) A two-step process

Fig. 8: Integrated biometric systems and spoofing countermeasures.

FLR reflects the percentage of spoofed trials misclassified as genuine trials whereas the FFR reflects the percentage of genuine trials misclassified as spoofed trials.

Even as independent sub-systems, spoofing countermeasures impact directly on the performance of the biometric system as a whole; while aiming to reduce vulnerabilities to spoofing, they also have potential to reject genuine trials [7]. Thus, while countermeasure sub-systems can be assessed independently, it is often of greater interest to assess the performance of the system as a whole, for example using a score
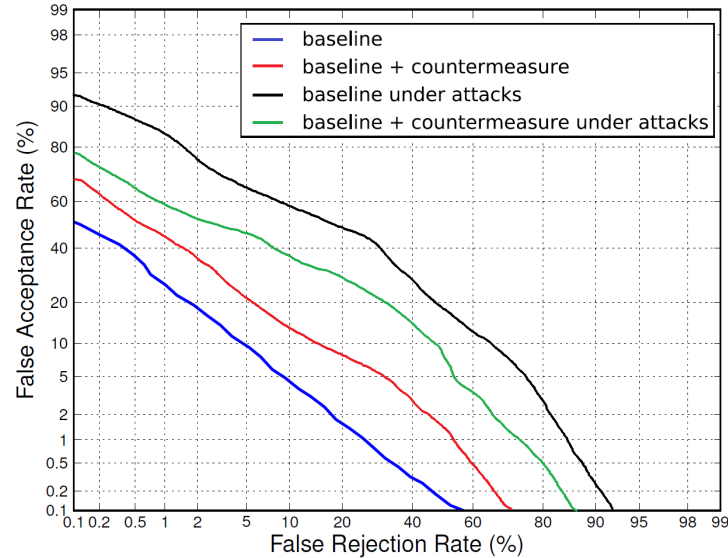
Fig. 9: The FAR and FRR of a biometric system with/without spoofing attacks and with/without countermeasures.

fusion strategy illustrated in Fig. 8(b).

Countermeasure sub-systems may alternatively be integrated, not in parallel, but in series with the biometric system as illustrated in Fig. 8(c). Assessment involves four different system configurations, for each of which a different DET profile is produced. Examples are illustrated in Fig. 9. The first configuration (lowest, blue profile) illustrates the performance of the baseline biometric system (no spoofing, no countermeasures). The second configuration (highest, black profile) illustrates the performance of the same system when subjected to spoofing attacks. The third configuration (second from highest, green profile) illustrates the improvement in performance with active countermeasures. While seemingly complete, a fourth configuration (second from bottom, red profile) is still needed to determine the performance of the integrated biometric and countermeasure systems in the absence of spoofing. This final configuration is essential in order to gauge the effect of the countermeasure on genuine trials which may be misclassified as spoofing attacks. To reiterate, all four configurations are needed to properly observe the complex impact of spoofing and countermeasures on integrated systems.

Even if it is the effect on overall recognition performance which is of greatest interest, there will always be an interest to first assess countermeasure performance independently. This configuration might be limited to system development. The second, fused approach is simple and straightforward. However, unless separate decisions are applied to the countermeasure and recognition sub-system scores, it does

not support the explicit detection of spoofed trials. The third approach is therefore the most appealing in practice, allowing for explicit spoof detection and an evaluation of countermeasure impacts on overall system performance.

Finally, whatever the approach to integration, performance is dependent on a separate countermeasure decision threshold τ^{CM}. Since it impacts upon overall performance, the FAR/FRR/SFAR can be determined for a range of different operating points specified, for example, by an application-dependent FFR. So as to reflect the likely performance for a variety of different application scenarios, the TABULA RASA project considered values of FFR$= \{1, 5, 10\}\%$. Even so, countermeasure integration and assessment is not a solved problem and very much a topic of ongoing research. This is discussed further in Section VII.

VI. Evaluation methodology: a case study in 2D face verification

As an illustration of the evaluation methodology in practice the following presents a case study in 2D face verification.

A. Previous work

The vast majority of past research in 2D face recognition has focused on maximising the discrimination between the faces of different persons [8]. Liveness detection has received considerably less attention, even though 2D face recognition systems have been known to be vulnerable to spoofing for some time.

Many face recognition systems can be spoofed by presenting to the camera photographs, videos or 3D masks of enrolled persons [2]. While make-up or plastic surgery are also viable spoofing attacks, photographs are among the most easily implemented, most effective and therefore the most likely in practice. Video attacks can be especially effective since they are dynamic signals which more closely resemble a genuine trial. Furthermore it is natural to assume that systems which are vulnerable to photo attacks will also be vulnerable to video attacks.

As is the case for the spectrum of other modalities, the typical countermeasure to prevent the spoofing of 2D face recognition systems involves liveness detection. Here, liveness indicators include eye blinking, changes in facial expression, mouth movements, estimates of skin texture, structure and motion analysis, and depth information etc. Multi-spectral and reflectance analysis has also been used successfully to differentiate between living faces and lifeless fakes [2]. Alternatively, face recognition can be combined with other biometric modalities such voice or gait recognition. A survey of face spoofing detection approaches can be found in [4].

B. **Baseline face biometric system**

The system combines a part-based face representation with Gaussian mixture models (GMMs) [9]. The system divides the face into blocks, and treats each block as a separate observation of the same underlying signal (the face). A feature vector is thus obtained from each block by applying the Discrete Cosine Transform (DCT). The distribution of the feature vectors is then modelled using GMMs.

For feature extraction, the face is normalised, registered and cropped. The cropped and normalised face is divided into blocks (parts) from each of which a feature vector is extracted. Each feature vector is treated as a separate observation of the same underlying signal and the distribution of the feature vectors is modelled using GMMs. The feature vectors from each block are obtained by applying the DCT. Once the feature vectors are calculated, feature distribution modelling is achieved by performing background model adaptation of GMMs. Background model adaptation first involves the training of a world (background) model Ω_{world} from a large dataset of faces. Client models Ω_{client}^{i} for client i are learned through the adaptation of the world model towards the observations of the client.

Verification is achieved by scoring an observation x against both the client (Ω_{client}^{i}) and world (Ω_{model}) models. They produce log-likelihood scores which are combined to give the single log-likelihood ratio (LLR). The LLR is used to assign the observation to the world class (not the client) or the client class based on a predefined threshold τ.

C. **Spoofing countermeasure**

Genuine and spoofed face images exhibit differences at the texture level. Spoofed face images are captured or acquired twice, first from the live client and second from the spoofing medium at recognition time. Accordingly spoofed face images tend to exhibit degraded facial texture due to reproduction or printing on the spoofing medium (a photograph) and the two imaging systems [10]. Accordingly, estimates of facial texture quality, here obtained using logistic regression [11], serve to identify spoofing attacks. The method is computationally efficient and does not require any additional user-cooperation; it is non-intrusive.

Fig. 10 illustrates a block diagram of the texture-based spoofing countermeasure. The approach analyses the texture of single facial images using two different texture descriptors: local binary patterns (LBP) which encodes the micro-texture patterns into a feature histogram, and features computed from the grey-level co-occurrence matrices (GLCM) which describe the distribution of different combinations of grey level pairs in an image block. The combination of these two measurements provides an effective representation of the overall facial texture quality. The facial texture descriptions are fed into logistic
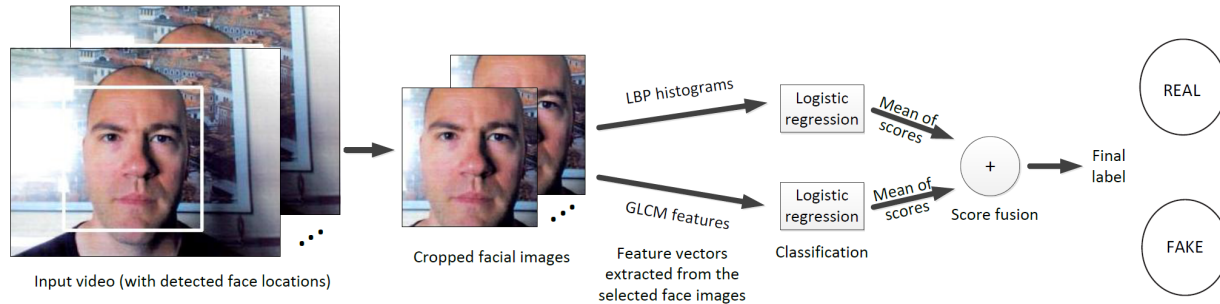
Fig. 10: A block diagram of the texture-based countermeasure.

regression classifiers. Score level fusion of the individual classifier outputs determine whether the facial image corresponds to a living person or a spoofed reproduction. As the database consists of video sequences, several cropped face images from each video are used for feature extraction at intervals of 0.6 seconds. The final score for the two individual texture representations is determined by averaging the scores of each face image. The outputs of each classifier are then fused using the sum rule with min-max score normalization [6] in order to obtain the final label for each video sequence.

D. **The face spoofing attack database**

Experiments were performed with the REPLAY-ATTACK face spoofing database [12]. It consists of 1300 samples comprising genuine trials, photo and video spoofing attacks. The database contains samples collected from 50 persons under varying lighting conditions. The data is split into 4 sub-groups comprising enrolment, training, development and test data. While the enrolment set includes samples collected from all clients, there is no client overlap in the training, development and test datasets.

All videos are generated by either having a (real) client trying to access a laptop through a built-in webcam or by displaying a photo or a video recording of the same client for at least 9 seconds. In total, 20 attack videos were recorded for each client whereas 6 videos were captured for real accesses. The enrolment set contains 100 videos (2 per client) and is used exclusively to evaluate baseline performance. The training set contains 60 real-accesses and 300 attacks. The development set contains 60 real-accesses and 300 attacks whereas the test set contains 80 real-accesses and 400 attacks. Examples of real accesses and spoofing attacks from the REPLAY-ATTACK database are shown in Fig. 11. A complete description of the database and associated protocol can be found in [12].

(a))   (b)

Fig. 11: The setup and sample images from the REPLAY-ATTACK database. Spoofed data collection using printed photographs (a) and examples of genuine accesses and spoofing attacks (b). Columns from left to right show examples of real accesses, printed photographs, mobile phone and tablet attacks.

E. Experimental results

Vulnerabilities to spoofing and countermeasure impacts were assessed according to the evaluation methodology introduced in this article. Results in the form of DET profiles are illustrated in Fig. 12. The first three figures (a), (b) and (c) represent overall system performance under spoofing attack for three different operating points where countermeasures are tuned to produce FFRs of 1%, 5% and 10%. For completeness a fourth figure, Fig. 12(d) is included to illustrate countermeasure performance in independence from biometric recognition. The EER of the countermeasures is in the order of 5%.

The texture-based countermeasure delivers significant improvements in robustness to spoofing at almost every operating point. For example, when the countermeasure is tuned to an FFR of between 1% and 5% as illustrated in Figs. 12(a) and (b), and for a fixed FRR of 1%, then the FAR under spoofing attack drops from 80% to just over 5%. However, if the countermeasure is tuned to a higher FFR such as 10% as illustrated in Fig. 12(c), then the FAR increases to approximately 20%. This is high-security configuration. Even if almost no spoofing trials are misclassified as genuine trials, the countermeasure misclassifies 10% of genuine trials as spoofed trials.

In this case, the countermeasure degrades usability; this effect is illustrated by the blue profile in Fig. 12(c). There is thus an inherent trade-off between the security and usability, which is heavily dependent on the modality, recognition system, countermeasure and database. For this particular configuration, a sensible compromise is achieved when the countermeasure FFR is tuned to 5%. Here, 5% of genuine trials will be misclassified as spoofed trials, whereas for all but the highest FRRs, only 5% of spoofing attacks
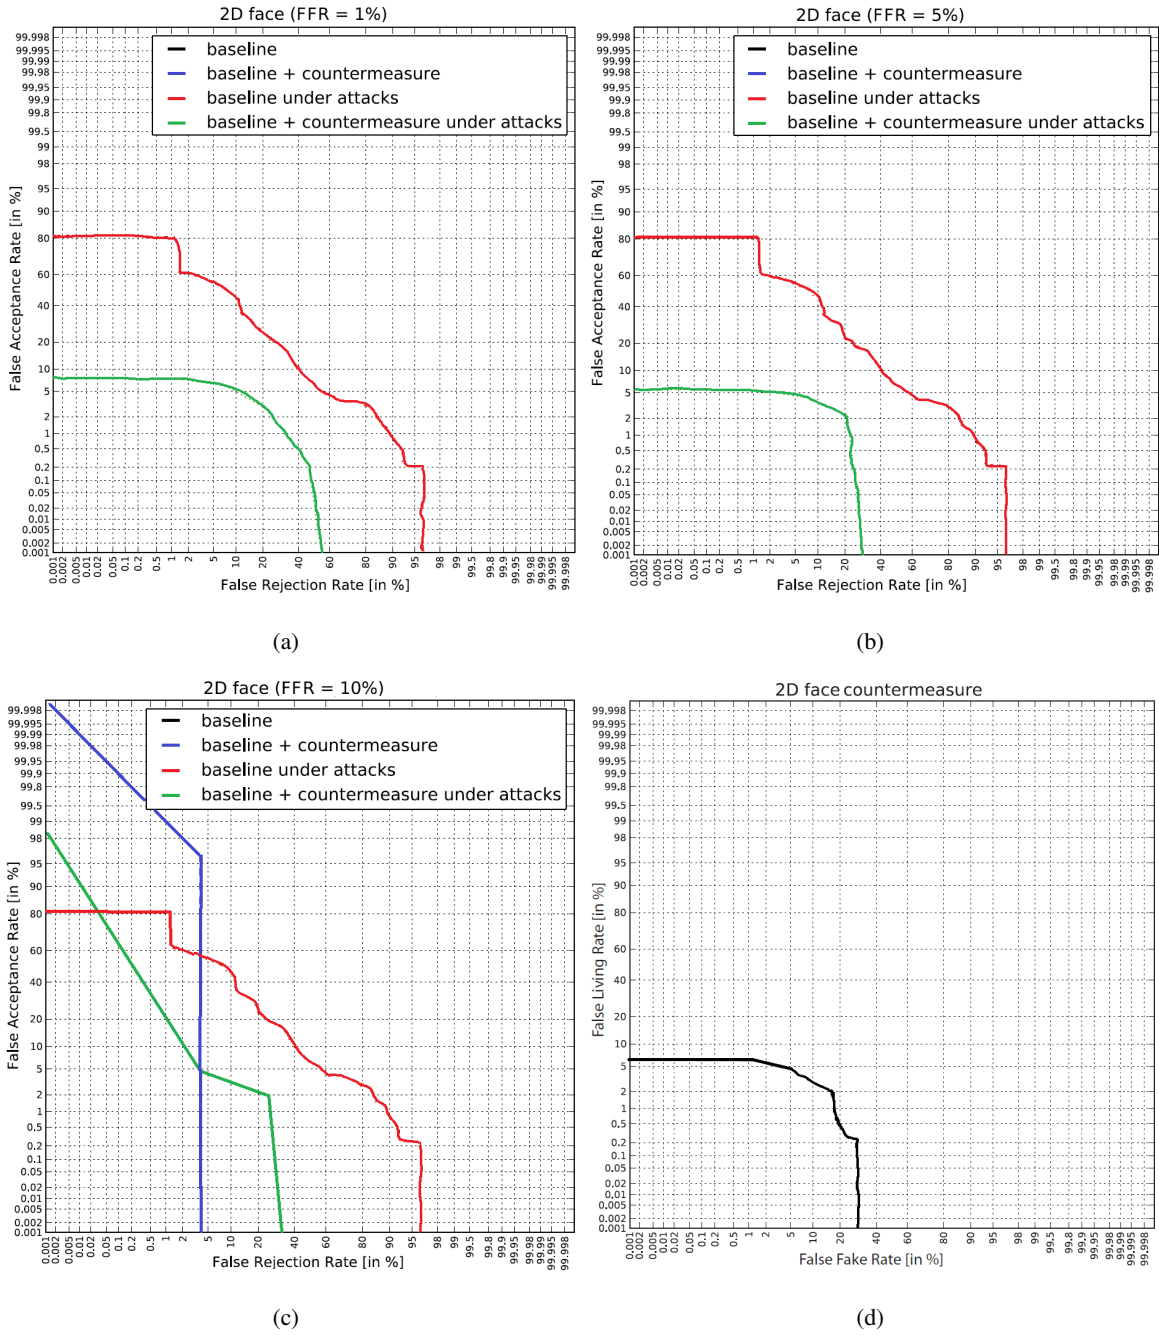
(a)

(b)

(c)

(d)

Fig. 12: DET profiles for the 2D face verification systems with/without spoofing and with/without countermeasures at different countermeasures operating points (FRRs of (a) 1%, (b) 5% and (c) 10%). Missing profiles indicate EER=FRR=FAR=0%. Also illustrated in (d) in a DET profile for the independent countermeasure.

will succeed. Between 40% and 80% of attacks would otherwise be successful without the texture-based countermeasure.

VII. LESSONS LEARNED

The evaluation methodology presented in this article is based upon work funded through the European TABULA RASA project. Through a collaboration involving a large team of researchers, the study included a wide range of biometric modalities including face (2D, 3D and multi-spectral), voice, gait, fingerprint, iris, vein and electro-physiological signals. The study established state-of-the-art authentication technologies for each modality, investigated the vulnerabilities of each to spoofing, and proposed novel countermeasures integrated and evaluated according to the methodology outlined in Sections IV and V. The findings described in the face recognition case study are illustrative of the general trend across all modalities considered in TABULA RASA. The following discusses some of the lessons learned and the most pressing directions for future research.

A. Vulnerabilities

Unless they are equipped with suitable countermeasures, all biometric systems were shown to be vulnerable to spoofing. Even so, some modalities (e.g. gait) are more robust than others (e.g. fingerprint), however, this should not be interpreted as meaning they are more reliable; in the absence of spoofing, fingerprint recognition generally outperforms gait recognition. Multi-modal biometric systems are also vulnerable and can be overcome by the spoofing of only a single modality [13].

B. Metrics and protocols

Spoofing and countermeasure assessment is considerably more complex than might first appear. Countermeasures tend to be developed and evaluated independently from recognition systems and although this methodology supports the comparison of different anti-spoofing approaches, it is the influence of countermeasures on overall system performance that is of greatest interest. New, standard metrics and protocols reflecting the robustness of integrated biometric and countermeasure systems should thus be adopted in the future.

C. Usability

While countermeasures are successful in reducing the vulnerability of biometric systems to spoofing, increased robustness comes at the expense of increased FRR. This impact on usability is generally

manageable; increases in the FRR are usually negligible in comparison to increases in the FAR caused by spoofing. It is stressed, however, that the trade-off is dependent on the modality and application.

D. Generalization

A large number of anti-spoofing studies have been reported in the literature and encouraging results have been reported for a small number of standard databases e.g. [14], [15], [16]. Spoofing attacks are, however, varying and unpredictable in nature and thus it is difficult to predict how countermeasures will generalise to spoofing attacks 'in the wild', where their true nature can never be known with certainty. Thus, as the field evolves, and as new forms of spoofing emerge, it will be necessary to collect databases with increasingly challenging and diverse spoofing trials, calling for new and increasingly generalised countermeasures. An alternative or complementary strategy involves one-class classification approaches [17]. As a form of anomaly detection which relies only on models of genuine data, they offer greater potential for generalisation.

E. Countermeasure fusion

The growing interest in spoofing and countermeasures and the number of diverse countermeasure strategies and algorithms in the literature lends support to research in fused countermeasures. While not increasing robustness to specific attacks, fused countermeasures offer a flexible anti-spoofing framework with which newly emerging vulnerabilities can be quickly patched with adaptable fusion strategies or the addition of new countermeasures.

F. Challenge-response countermeasures

Interactive, challenge-response countermeasures require a user to perform a randomised, specific action such as smiling or blinking. The correct response to the given challenge infers liveness. While potentially more intrusive, challenge-response countermeasures can be complementary to the more traditional approaches to liveness detection. Further work is required to validate their potential.

G. Combined hardware and software-based countermeasures

The majority of previous work, including all that within TABULA RASA, has investigated software-based countermeasures. Hardware-based countermeasures and new sensors have great potential to detect more elaborate spoofing attacks and thus a new generation of countermeasures which combine hardware and software-based approaches may be needed.

VIII. Outlook

Until relatively recently, the main focus of biometrics research has centred on the discrimination between genuine and zero-effort impostor trials. In the face of well-acknowledged vulnerabilities, greater effort to develop countermeasures will be required in the future in order to defend against concerted-effort spoofing attacks. The growing body of related literature and number of competitive evaluations is testament to the increasing interest and importance of this research.

The past work is characterised by the lack of a standard evaluation methodology which is needed to assess the influence of countermeasures on biometric system performance. This article presents what is, to the best of the authors' knowledge, the first proposal for a formal evaluation methodology. Needless to say, however, further work is required to extend and adapt the methodology in view of the lessons learned through recent work.

While the performance of spoofing countermeasures proposed thus far give cause for optimism, their generalisation to previously unseen spoofing attacks remains unclear. This leads to a number of research directions for the future, including networks of independent, fused countermeasures and one-class classification strategies. As the field evolves, new and more challenging databases will be essential in order for biometric system developers to stay one step ahead of the fraudsters. While extremely difficult, it will be critical to estimate the reliability of countermeasures in practical application scenarios including, not just their ability to detect spoofing, but also their impact on system usability.

References

[1] A. Jain, A. Ross, and S. Pankati, "Biometrics: A tool for information security," *IEEE Transactions on Information Forensics and Security*, vol. 1, no. 2, pp. 125–143, June 2006.

[2] S. Marcel, M. Nixon, and S. Z. Li, *Handbook of Biometric Anti-Spoofing*. Springer, 2014.

[3] E. Marasco and A. Ross, "A survey on antispoofing schemes for fingerprint recognition systems," *ACM Comput. Surv.*, vol. 47, no. 2, pp. 28:1–28:36, Nov. 2014.

[4] J. Galbally, S. Marcel, and J. Fierrez, "Biometric anti-spoofing methods: A survey in face recognition," *IEEE Access*, vol. 2, pp. 1–23, December 2014.

[5] Z. Wu, N. Evans, T. Kinnunen, J. Yamagishi, F. Alegre, and H. Li, "Spoofing and countermeasures for speaker verification: A survey," *Speech Communication*, vol. 66, no. 0, pp. 130 – 153, 2015.

[6] J. Fierrez, "Adapted fusion schemes for multimodal biometric authentication," Ph.D. dissertation, Universidad Politecnica de Madrid, May 2006.

[7] I. Chingovska, A. Anjos, and S. Marcel, "Biometrics evaluation under spoofing attacks," *IEEE Transactions on Information Forensics and Security*, vol. 9, no. 12, pp. 2264–2276, December 2014.

[8] S. Z. Li and A. K. Jain, Eds., *Handbook of Face Recognition*. New York: Springer, 2nd edition, 2011.

[9] L. El-Shafey, C. McCool, R. Wallace, and S. Marcel, "A scalable formulation of probabilistic linear discriminant analysis: Applied to face recognition," *IEEE Trans. on Pattern Anal. and Machine Intel.*, vol. 35, no. 7, pp. 1788–1894, July 2013.

[10] J. Galbally, S. Marcel, and J. Fierrez, "Image quality assessment for fake biometric detection: Application to iris, fingerprint and face recognition," *IEEE Trans. on Image Processing*, vol. 23, no. 2, pp. 710–724, February 2014.

[11] J. Maatta, A. Hadid, and M. Pietikainen, "Face spoofing detection from single images using texture and local shape analysis," *IET Biometrics*, vol. 1, no. 1, pp. 3–10, 2012.

[12] T. de Freitas Pereira, J. Komulainen, A. Anjos, J. M. D. Martino, A. Hadid, M. Pietikainen, and S. Marcel, "Face liveness detection using dynamic texture," *EURASIP Journal on Image and Video Processing*, vol. 2014:2, January 2014.

[13] G. L. Marcialis, G. Fumera, and B. Biggio, *Encyclopedia of Biometrics*. Springer, 2014, ch. Anti-spoofing: Multimodal.

[14] L. Ghiani, D. Yambay, V. Mura, S. Tocco, G. L. Marcialis, F. Roli, and S. Schuckers, "Livdet 2013 – fingerprint liveness detection competition," in *Proc. IEEE/IAPR Int. Conf. on Biometrics*. IEEE Press, June 2013.

[15] I. Chingovska, J. Yang, Z. Lei, D. Yi, S. Z. Li, O. Kahm, C. Glaser, N. Damer, A. Kuijper, A. Nouak, J. Komulainen, T. Pereira, S. Gupta, S. Khandelwal, S. Bansal, A. Rai, T. Krishna, D. Goyal, M.-A. Waris, H. Zhang, I. Ahmad, S. Kiranyaz, M. Gabbouj, R. Tronci, M. Pili, N. Sirena, F. Roli, J. Galbally, J. Fierrez, A. Pinto, H. Pedrini, W. S. Schwartz, A. Rocha, A. Anjos, and S. Marcel, "The 2nd competition on counter measures to 2d face spoofing attacks," in *Proc. IEEE/IAPR Int. Conf. on Biometrics*. IEEE Press, June 2013.

[16] Z. Wu, T. Kinnunen, N. Evans, J. Yamagishi, C. Hanilci, M. Sahidullah, and A. Sizov, "ASVspoof 2015: the first automatic speaker verification spoofing and countermeasures challenge," in *INTERSPEECH*, 2015.

[17] F. Alegre, A. Amehraye, and N. Evans, "A one-class classification approach to generalised speaker verification spoofing countermeasures using local binary patterns," in *Proc. IEEE 6th Int. Conf. on Biometrics: Theory, Applications and Systems (BTAS 2013)*, 2013, pp. 1–8.

Abdenour Hadid received the Doctor of Science in Technology degree in Electrical and Information Engineering from the University of Oulu, Finland, in 2005. Since 2010, he is an Adjunct Professor and Academy Research Fellow in the Center for Machine Vision Research of the University of Oulu. His research interests include computer vision and pattern recognition and particularly face analysis and biometrics. He made significant contributions to the state-of-the-art and his work is gaining increasing interest in the scientific community. According to Google scholar (May 2015), his h-index is 24 and his work is already cited more than 5600 times.

Nicholas Evans is an Assistant Professor at EURECOM, France. He received the MEng (1999) and PhD (2003) degrees from the University of Wales, Swansea, UK. He is the lead guest editor for the IEEE T-IFS Special Issue on Biometric Spoofing and Countermeasures and the IEEE SPM Special Issue on Biometric Security and Privacy. He is an Associate Editor of the EURASIP Journal on Audio, Speech and Music Processing and is a member of the IEEE Speech and Language Technical Committee.

Sébastien Marcel is researcher at the Idiap Research Institute, Switzerland, where he heads the biometrics group. He is also lecturer at the Ecole Polytechnique Fédéral de Lausanne. He is Associate Editor of IEEE Transactions on Information Forensics and Security (T-IFS), co-Editor of the 'Handbook of Biometric Anti-Spoofing', Guest Editor of the IEEE T-IFS Special Issue on Biometric Spoofing and Countermeasures and of the IEEE SPM Special Issue on Biometric Security and Privacy.

Julian Fierrez received the MSc and PhD degrees from UPM in Madrid, Spain, in 2001 and 2006, respectively. Since 2004 he has been affiliated with Universidad Autonoma de Madrid, where he is currently an Associate Professor. From 2007 to 2009 he was a Marie Curie Postdoc at Michigan State University, USA. His research interests include image processing, pattern recognition, authentication using biometrics such as fingerprints and handwritten signatures, and security of person authentication systems. He is actively involved in multiple EU projects, and has received multiple research distinctions including the EURASIP Best PhD Award in 2012.