Pattern Recognition ■ (■■■■) ■■■–■■■

Contents lists available at ScienceDirect

Pattern Recognition

journal homepage: www.elsevier.com/locate/pr

Robust multimodal face and fingerprint fusion in the presence of spoofing attacks

Peter Wild *, Petru Radu, Lulu Chen, James Ferryman

University of Reading, School of Systems Engineering, Whiteknights, Reading RG6 6AY, United Kingdom

ARTICLE INFO

Article history:
Received 28 February 2015
Received in revised form
15 June 2015
Accepted 7 August 2015

Keywords:
Multibiometrics
Anti-spoofing
1-Median

ABSTRACT

Anti-spoofing is attracting growing interest in biometrics, considering the variety of fake materials and new means to attack biometric recognition systems. New unseen materials continuously challenge state-of-the-art spoofing detectors, suggesting for additional systematic approaches to target anti-spoofing. By incorporating liveness scores into the biometric fusion process, recognition accuracy can be enhanced, but traditional sum-rule based fusion algorithms are known to be highly sensitive to single spoofed instances. This paper investigates 1-median filtering as a spoofing-resistant generalised alternative to the sum-rule targeting the problem of partial multibiometric spoofing where m out of n biometric sources to be combined are attacked. Augmenting previous work, this paper investigates the dynamic detection and rejection of liveness-recognition pair outliers for spoofed samples in true multi-modal configuration with its inherent challenge of normalisation. As a further contribution, bootstrap aggregating (bagging) classifiers for fingerprint spoof-detection algorithm is presented. Experiments on the latest face video databases (Idiap Replay-Attack Database and CASIA Face Anti-Spoofing Database) and fingerprint spoofing database (Fingerprint Liveness Detection Competition 2013) illustrate the efficiency of proposed techniques.

1. Introduction

Fingerprint and face biometrics as most widely adopted traits are being exposed to an increasing threat of presentation attacks. Consequently, there are numerous studies [1,2] and open challenges [3,4] on anti-spoofing techniques assessing the spoofing detector's ability to distinguish between genuine and fake attempts for especially these two traits. Recently, the integration of anti-spoofing scores with recognition scores has received considerable attention [5–7]. The standard approach, as outlined in [5], has been to reject spoofed samples before comparing them against the gallery template. However, recognition scores can be helpful in the probe-attack spoofing detection problem and liveness scores can impact on the recognition task. Considering imposters with access to fake fingers or face photographs reveals an impact on overall accuracy (shifted imposter score distribution for non-zero-effort attempts [7]) and assuming a correlation between successful spoofs achieving a higher score and their corresponding liveness score is likely (and shown) to help in the final judgment of the decision task, especially in an ensemble of classifiers where this paper looks for outliers. It is therefore useful to investigate the benefits of dealing with a holistic (liveness and verification) multi-class problem rather than two separate classification problems (live vs. fake and genuine vs. impostor). If a system involves multiple modalities there is an even larger variety of different ways to treat the problem of combining liveness and recognition scores. Multibiometrics using face and fingerprint biometrics comes with many benefits including expected increased accuracy, higher universality (absence of single characteristics), efficiency (fast indexing), but its robustness to spoofing attempts has been shown to be compromised [8,9]. Furthermore, with the inclusion of multiple modalities the attacker has an even more extended choice to select the easiest modality to be attacked. It is therefore desirable to find new techniques coping with spoofing attacks, which are subject to investigation in this paper. The paper focuses on three objectives: (1) investigation of spoofing robustness in multibiometrics; (2) development of novel methods towards anomaly detection for increased systematic anti-spoofing; and (3) proposition of a novel bootstrap aggregating (bagging) of classifiers method combining features in fingerprint counter-spoofing.

With regard to the first topic on spoofing robustness in multibiometrics, the paper tests degradation in accuracy for the "partial multibiometric spoofing" scenario, where m out of n samples are spoofed, highlighting the tradeoff between accuracy and security for different fusion methods. Fig. 1 illustrates this concept. The

* Corresponding author. Tel.: +44 118 378 7633; fax: +44 118 975 1994.
E-mail addresses: p.wild@reading.ac.uk (P. Wild), p.radu@reading.ac.uk (P. Radu), l.chen@reading.ac.uk (L. Chen), j.m.ferryman@reading.ac.uk (J. Ferryman).

http://dx.doi.org/10.1016/j.patcog.2015.08.007

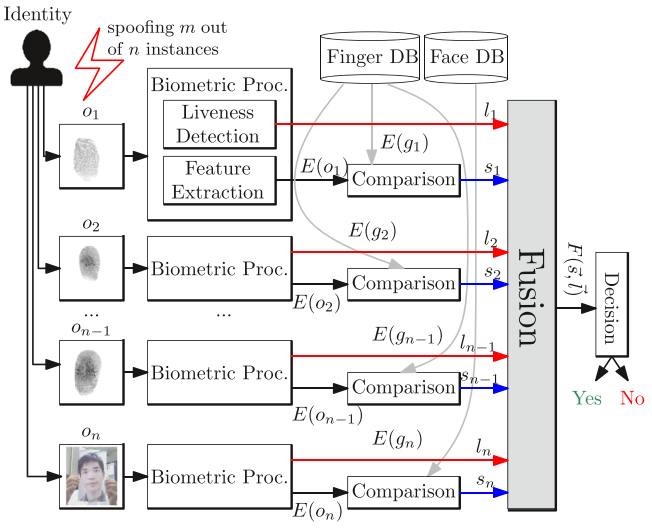

Fig. 1. Partial multibiometric spoofing of observations o_i given templates g_i fusing scores s_i and liveness values l_i.

sensitivity of a recognition and liveness fusion method with regard to spoofing is especially interesting in multimodal configuration, where scores originate from different underlying distributions and multiple traits facilitate a selection of the modality to be attacked. The paper analyses the impact of the number of spoofed fingers or spoofed face on accuracy using the latest biometric datasets. The relative robustness of several score-level fusion rules can be used to choose the most robust fusion rule [9].

As a second outlined contribution, this paper presents a novel multibiometric spoofing-aware fusion method following the idea of anomaly detection and extending research in [10] to multiple modalities. This paper investigates 1-median-based fusion using outlier detection applied in a multibiometric setup. Note that the extension to multiple modalities raises further questions with regard to normalisation. For different modalities, scores generally follow different distributions. Therefore, counter-spoofing is much more challenging than for single-modality approaches, including multi-instance or multi-algorithm approaches. Further, this work presents further theoretical considerations and discusses parameter choice in detail. Recognition scores and liveness scores are likely to be dependent, as spoofing tries to achieve a high recognition score in order to successfully claim the alien (spoofed) identity. In partial multibiometric spoofing this information can be used to further discriminate between genuines and impostors. Despite spoofing sensitivity of traditional fusion techniques, it is a reasonable assumption to claim a higher difficulty for attackers to spoof multiple modalities at the same time or even to obtain the necessary samples to produce a fake fingerprint or face mask. On the other hand, special spoofing-robust fusion schemes might exhibit a reduced level of accuracy. This trade-off between cost and security to limit drawbacks [5,11] is investigated.

Third and last, as a by-product of evaluations the paper further presents a novel spoofing detector again employing a fusion principle: bootstrap aggregating (bagging) of classifiers. This technique is employed in combining the decision outcome of multiple different classifiers. Using also multiple features to be more robust vs. changes in materials (see [12]), the paper aims at investigating this technique in the employed system as an anti-fingerprint spoofing technique towards integral fusion concepts in robust anti-spoofing. Bagging is shown to outperform state-of-the-art detectors on the most challenging LivDet 2013 Crossmatch subset database.

The remainder of this paper is organised as follows: Section 2 introduces the problems of anti-spoofing and spoofing-aware fusion in biometrics. The proposed methods of bagging for spoof-detection and 1-median filtering for spoofing-resistant multibiometric fusion are outlined in Section 3. Section 4 highlights experimental results with regard to the proposed and investigated techniques. This includes a discussion of methods towards anomaly detection in multibiometrics, highlighting parameter choice and optimisation for the proposed 1-median filtering. Section 5 concludes this paper with an outlook on future work.

2. Related work

There are several anti-spoofing or liveness detection algorithms extracting features (usually trained for modality, sensor, material, etc.), in order to determine whether a biometric sample is either *live* or *fake*. For evaluation purposes, *ferrlive* (rate of misclassified live samples) and *ferrfake* (rate of misclassified fake samples) are employed. Whereas for individual modalities the anti-spoofing problem is well defined and evaluated separately from biometric system performance, research on fusion between match scores and liveness factors is still in its infancy [13]. Recently, [14] suggested a framework for verification systems under spoofing attacks. Within the framework [8] adopted in this paper, liveness and recognition scores are combined considering the scenario of probe-spoofing only (i.e. no gallery-spoofing, enforced by e.g., attended enrolment). Formally, given a vector of biometric observation (units, e.g. fingers, eyes) $\vec{o} = (o_1, \ldots o_n)$ from one or more modalities, and corresponding claimed identity template $\vec{g} = (g_1, \ldots, g_n)$, the task of the fusion module F is to compute a unified decision score, using comparison scores $\vec{s} = (s_1, \ldots, s_n)$ and (probe) liveness values $\vec{l} = (l_1, \ldots, l_n)$, so that the verification task V (authentication based on threshold η) can be formulated as follows:

$$V(\vec{o}, \vec{g}) := \begin{cases} \text{accept} & \text{if } F(\vec{s}, \vec{l}) \geq \eta; \\ \text{reject} & \text{else.} \end{cases} \quad (1)$$

Let i be the current index and $E(o_i), E(g_i)$ refer to extracted (modality-specific) features of samples, then $s_i = C(E(o_i), E(g_i)) \in [0, 1]$ is used to denote the normalised comparison result of o_i, g_i and $l_i = L(o_i) \in [0, 1]$ denote the likeliness of a genuine (live) sample. Clearly, it is desirable to find a method F unaffected in performance if m out of the n elements of $\vec{o}$ are spoofed. This testing setup is referred to as "partial multibiometric spoofing", introduced in [10] and extended in this work towards multiple modalities. Note that this notion of live or spoofed probes vs. always-live enrolled gallery samples (assuming attended enrollment) leads to a simpler modelling (2 classes distinguishing live probe from spoof or live, but different sources) than in the general asymmetric case (8 classes based on live/spoof probe, live/spoof gallery sample, and same/different source) or symmetric case (6 classes) [7], fully concentrating on a dichotomous authentication task, which can be evaluated in the traditional way using receiver operating characteristics.

2.1. On combining anti-spoofing and recognition

Marasco et al. [5] are among the first considering fusion of liveness with recognition scores separately for each modality, using simple rejection of spoofed samples. If a spoofing attempt is indicated, the current modality matching score is ignored. This initial study is extended in [15] evaluating sequential fusion, classifier fusion, and Bayesian Belief Networks for combining match scores and liveness measures, highlighting the superiority of the latter method for the LivDet2009 dataset but also that accuracy is decreased when taking liveness detection into account. Chingovska et al. [6] evaluate binary decision rules and Logistic Regression (LR) as decision and score-level fusion techniques combining face recognition and liveness scores addressing the integration (but neglecting the partial spoofing problem) of liveness. They report higher resistance to spoofing attacks (91.54% vs. 10%) but are outperformed by

Please cite this article as: P. Wild, et al., Robust multimodal face and fingerprint fusion in the presence of spoofing attacks, Pattern Recognition (2015), http://dx.doi.org/10.1016/j.patcog.2015.08.007