

Battista Biggio, Giorgio Fumera, Paolo Russu, Luca Didaci, and Fabio Roli

Adversarial Biometric Recognition

A review on biometric system security from the adversarial machine-learning perspective



Biometrics Security and Privacy Protection

In this article, we review previous work on biometric security under a recent framework proposed in the field of adversarial machine learning. This allows us to highlight novel insights on the security of biometric systems when operating in the presence of intelligent and adaptive attackers that manipulate data to compromise normal system operation. We show how this framework enables the categorization of known and novel vulnerabilities of biometric recognition systems, along with the corresponding

attacks, countermeasures, and defense mechanisms. We report two application examples, respectively showing how to fabricate a more effective face spoofing attack, and how to counter an attack that exploits an unknown vulnerability of an adaptive face-recognition system to compromise its face templates.

INTRODUCTION

Adversarial machine learning is a novel research field that was born in response to the increasing use of pattern recognition and machine-learning techniques, including signal processing ones, in security-related applications such as biometric identity recognition,

spam, and malware detection. In these applications, intelligent and adaptive adversaries are interested in subverting system operation; e.g., nonauthorized users may aim to gain access to a resource secured by a biometric identity recognition system. Despite the fact that pattern recognition and machine-learning algorithms have enabled the development of more effective recognition systems, they have not been originally designed to operate in adversarial settings. In particular, their underlying assumption of data stationarity (i.e., that training and testing data follow the same distribution) is likely to be violated in adversarial environments. As a consequence, these algorithms can introduce additional, specific vulnerabilities that can be exploited by carefully crafted attacks to cause different security violations, including denial of service and missed detection of intrusive attempts. This may eventually compromise the whole system security. Even if countermeasures and novel algorithms have been proposed to improve security against these sophisticated attacks, they will not stop adversaries from developing novel ways of misleading such defense systems, engendering a long-lasting arms race.

To date, research efforts in adversarial machine learning have focused on identifying different kinds of potential attacks against machine-learning and pattern recognition algorithms, and developing the corresponding countermeasures to improve robustness in adversarial settings [1]–[4]. According to the security-by-design paradigm, ongoing work is also addressing the issue of extending learning theory and methods to explicitly account for the presence of malicious adversaries that can undermine algorithm operation; e.g., in [2] and [3], traditional performance evaluation methods have been extended to allow for a systematic evaluation of the security of pattern classifiers in adversarial settings, providing a better understanding of the system performance both in the absence and in the presence of well-crafted attacks. The relevance of these issues is also witnessed by an increasing number of publications and events, e.g., the Neural Information Processing Systems Workshop on Machine Learning in Adversarial Environments for Computer Security [5] and the more recent Dagstuhl Perspectives Workshop on Machine-Learning Methods for Computer Security [6].

Biometric identity recognition is a clear example of a widespread and still growing application field in which security is a key issue and pattern recognition techniques play a major role. Different vulnerabilities of biometric systems, specific attacks that can exploit them, and corresponding countermeasures have been analyzed in the literature [7], [8] and in research projects. For instance, the recent EU FP7 Tabula Rasa project has carried out an extensive analysis on spoofing attacks (i.e., attacks involving the submission of a fake biometric trait, like a gummy finger, to impersonate an authorized user), and on the development of possible countermeasures, like liveness detection techniques. Moreover, several approaches for the analysis and assessment of biometric systems security have been proposed; see, e.g., the ISO/IEC 19792:2009 implementation specifics for the security evaluation of biometrics, and the National Institute of Standards and Technology Common Criteria for Information Technology Security Evaluation. However, all existing efforts disregarded the potential, specific vulnerabilities introduced by pattern recognition algorithms used in biometric systems, and thus the investigation of the corresponding

attacks and countermeasures. We argue that looking at biometric system security from the perspective of adversarial machine learning not only provides an original categorization of existing attacks against such systems, but it also allows us to consider more sophisticated attacks targeting vulnerabilities of the learning algorithms used in these systems, along with the countermeasures already proposed in the field of adversarial machine learning.

Based on the above motivations, in the following we first provide a concise overview of adversarial machine learning, to introduce kindly the readers to this recent research field; we use popular attacks against biometric systems, such as spoofing attacks, as running examples to make our explanation clearer. We then review the security of biometric identity recognition systems by showing how recent theoretical results and systematization efforts from this field enable: 1) the definition of a more complete taxonomy of attacks against biometric systems, based on a formal attacker's model explicitly accounting for her knowledge and capability, which allows one to identify novel attack scenarios associated to specific vulnerabilities of machine-learning and pattern recognition algorithms, besides encompassing known attacks; and 2) the design of the corresponding countermeasures, building on solutions proposed in adversarial machine learning, which can give rise to the design of novel, secure-by-design algorithms capable of improving adversarial biometric identity recognition. We finally discuss two application examples of the possible aforementioned achievements. We first show how a skilled attacker may fabricate more effective face-spoofing attacks, and then highlight a new vulnerability of adaptive biometric systems, devising the corresponding attack and a possible countermeasure.

The main goal of this article is to provide the readers of this magazine, and researchers in biometrics, a gentle introduction to adversarial machine learning, and a well-structured review of the state of the art on biometric security in light of the most recent findings in the area of adversarial machine learning.

ADVERSARIAL MACHINE LEARNING: AN OVERVIEW

During the last several decades, the increasing variability and sophistication of attack threats, in response to the growing complexity and amount of vulnerable attack points in security systems, has favored the adoption of machine-learning and pattern recognition techniques to timely detect variants of known and never-before-seen attacks. These techniques can, however, exhibit intrinsic vulnerabilities that can be exploited by skilled attackers, perpetuating their arms race against system designers. Adversarial machine learning aims at countering this phenomenon by focusing on vulnerabilities of learning algorithms. It attempts to anticipate the adversary's strategy by identifying novel threats and devising the corresponding countermeasures before system deployment. In practice, it follows a proactive rather than a reactive approach. The first step toward the aforementioned goal has been the proposal of a taxonomy categorizing attacks against learning algorithms along three axes [1], [2]: 1) the attack influence, which can be exploratory, if the adversary can only manipulate the testing data, or causative, if she can modify also the training data; 2) the attack specificity, which ranges from targeted to indiscriminate, depending on whether the classification of a set of specific samples

or any of them is affected by the attack; 3) the security violation, which can be an integrity violation, if the adversary is allowed to access a restricted service or resource (e.g., an impostor gaining access to a genuine client's account [3], [9]); an availability violation, if legitimate users are denied access or normal system operation is compromised (e.g., misclassifying legitimate e-mails as spam); and a privacy violation, if the adversary is able to exploit confidential information about the system (e.g., the clients' templates in a biometric recognition system [10]–[12]).

To date, several vulnerabilities and attacks against different learning algorithms (e.g., support vector machines and neural networks) have been investigated, along with the proposal of possible countermeasures [13]–[15]. We will summarize the main existing countermeasures in the remainder of this article, discussing how they can be exploited to improve the security of machine-learning algorithms used in biometric systems. In particular, to anticipate the adversary's strategy, the existing work in adversarial learning simulates attacks, based on more or less explicit models of the adversary. We recently formalized this approach within a general framework, proposing a formal model to characterize the adversary's behavior [3], [4]. We summarize our model next and exploit it in the remainder of this article to characterize and understand security of biometric systems under an adversarial machine-learning perspective.

Our model generalizes and encompasses other models proposed in the area of adversarial machine learning [1], [2], making explicit assumptions on the attacker's goal, knowledge of the targeted system, and capabilities of manipulating the input data or the system's components.

The goal has to be defined according to the desired security violation and attack specificity. Knowledge of system components (e.g., the kind of decision function and its parameters, or how a component operates) can be perfect or limited, and feedback on the classifier's decisions can also be exploited [13], [16], [17]. The capability is defined as the attack influence, based on how the adversary can affect training and testing data (e.g., which features can be manipulated and how, according to application-specific

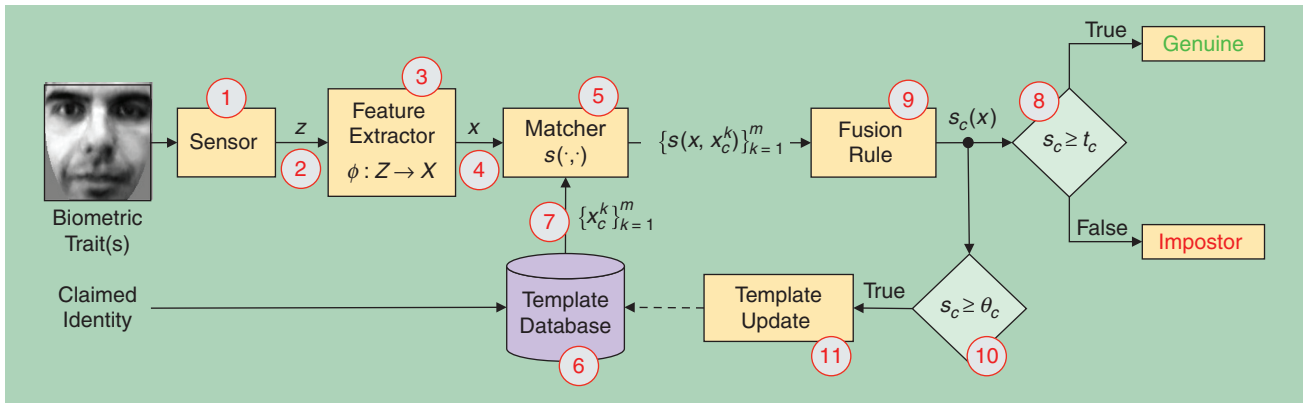
constraints). An attack strategy can be then defined based on the previous elements, to implement the attack. In formal terms, we assume the attacker's goal to be quantified by a function $g(a, \theta)$ measuring the extent to which an attack strategy a from a feasible strategy set \mathcal{A} fulfills the attacker's goal. The feasible set \mathcal{A} has to be defined according to the attacker's capability of manipulating the input data and system components, while the attacker's knowledge is encoded by the parameter vector $\theta \in \Theta$. Under this setting, the optimal attack strategy corresponds to the solution of the following optimization problem:

$$\max_{a \in \mathcal{A}} g(a; \theta). \quad (1)$$

Although this formulation may seem rather abstract at this stage, it enables us to consider trivial and sophisticated attacks under a consistent view, as shown in the remainder of this article.

BIOMETRIC RECOGNITION SYSTEMS UNDER ATTACK

Biometric recognition systems operate either in enrollment or in recognition mode [8]. During enrollment, each client provides his biometric traits and identity, in the presence of a human supervisor. A set of reference templates for each client is then stored in the template database along with the corresponding identity. During recognition, the biometric system is expected to recognize a previously enrolled client by comparing the submitted traits with those stored in the template database. Biometric systems may operate either in a verification or in an identification setting. In verification settings, biometric systems are often used to control access to protected resources, including confidential information or services. A user aiming to access them has to provide his/her biometric trait and claim an identity. The system then verifies whether the claim is genuine (i.e., the user's identity is the claimed one) or not (i.e., the user is an impostor trying to impersonate another client), and allows access only in the former case. This procedure is illustrated in Figure 1, which is general enough to also account for multibiometric systems; in this case, the fusion



[FIG1] The architecture of a biometric verification system and corresponding attack points, highlighted with red-circled numbers. During verification, the image $z \in \mathcal{Z}$ (e.g., a face image) acquired by the sensor is processed by a feature extractor $\phi: \mathcal{Z} \rightarrow \mathcal{X}$ to obtain a compact representation $x \in \mathcal{X}$ (e.g., a graph). The templates $\{x_c^k\}_{k=1}^m$ of the claimed identity c are retrieved from the template database, and compared to x using a matching algorithm $s: \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$. The resulting scores $\{s(x, x_c^k)\}_{k=1}^m$ are combined by a fusion rule, producing an aggregated score $s_c(x)$ that expresses the degree to which x is likely to belong to c . The score $s_c(x)$ is then compared with a decision threshold t_c to decide whether the claim is genuine or impostor. If a template self-update is implemented, and $s_c(x)$ is not lower than a self-update threshold θ_c , one of the templates in $\{x_c^k\}_{k=1}^m$ is updated depending on x , according to a given policy.

rule s_c should aggregate the matching scores coming from all biometric traits. In identification settings, instead, no identity claim is made: a user provides only the requested biometric trait x , and the system is expected to correctly recognize the corresponding identity among those in the template database, by matching x against all the known clients' templates; the corresponding scores $s_c(x)$ are then sorted in descending order to provide a list of the most likely candidate identities.

To account for natural changes of biometric traits over time (i.e., biometric aging), and changes in the environmental or acquisition conditions during verification, adaptive biometric systems have been proposed. They enable an update of the stored templates automatically during verification [18], [19]. A popular technique is template self-update: if the submitted trait is sufficiently similar to the reference templates of the claimed identity, one of such templates is updated by exploiting information coming from the submitted trait, according to a given policy. A simple update policy is called *nearest-out self-update*, as it replaces the most similar template to the submitted trait with the latter [20], [21].

THE ATTACK SURFACE

Previous work has identified the main attack points and vulnerabilities of biometric recognition systems, along with the corresponding attacks [7], [8]. First, any system is subject to intrinsic failures not produced by adversarial attempts, i.e., rejected genuine claims and accepted zero-effort impostors (i.e., impostors that do not exert any special effort to intrude). Besides this, a number of adversarial attacks have been also considered in early work, leading to the identification of eight potentially vulnerable attack points highlighted by the red-circled numbers 1–8 in Figure 1 [7]. Additionally, we consider points 9–11: they correspond to vulnerabilities of adaptive biometric systems that update clients' templates during operation, which we recently exploited to implement a template poisoning attack [20], [21]. The set of all attack points defines the attack surface of a biometric recognition system. The corresponding attacks can be categorized into four main groups according to the targeted system component [8]: 1) attacks to the sensor (point 1), 2) to interfaces and channels connecting different modules (points 2, 4, and 7), 3) to processing modules and algorithms (points 3, 5, and 8–11), and 4) to the template database (point 6). We discuss them below, along with the corresponding countermeasures proposed so far, which are also summarized in Table 1. It is also worth mentioning here a special category of attacks, known as *insider attacks*, where the attacker is

colluded with a system administrator or exercises coercion to escalate privileges [8].

Spoofing attacks consist of fabricating a fake biometric trait to impersonate an enrolled client. They target the sensor (point 1), so they are also referred to as direct attacks. Current defenses are based on liveness detection methods, which aim to verify whether the submitted trait is “alive” or “fake” by looking at specific patterns (e.g., perspiration patterns during fingerprint acquisition, or eye blinking during face verification). Multibiometric systems have been also proposed as a defense; however, to avoid spoofing them by only using a single fake trait, the matching scores coming from the different traits should be properly combined, using a secure score-level fusion rule [9], [22].

Replay attacks can be staged at interfaces between modules by replaying a stolen image of the biometric trait of the targeted client to the feature extractor (point 2), or directly the corresponding feature values to the matcher (point 4). An attacker may even replay a signal to replace the features of a given template of the claimed identity (point 7). This attack can be clearly staged if the corresponding communication channels are insecure, but also over encrypted channels, as the encrypted signal can be stolen and replayed into the channel directly. This can be avoided by encrypting a time stamp into the signal, or using challenge-response mechanisms. Another possible countermeasure is physical isolation, to avoid sending data over insecure channels (e.g., the Internet) subject to man-in-the-middle attacks. A popular example of physical isolation is the use of smart cards performing match-on-card operations. However, this technique has its own disadvantages, including limitations in terms of computational resources and memory, and the fact that the user should always use the smart card to be authenticated [8].

Hill-climbing attacks, similarly to replay ones, affect insecure communication channels between modules, and, in particular, points 2 and 4 in Figure 1. Their goal is to reconstruct a template image by iteratively sending a bunch of slightly perturbed images to the feature extractor (point 2), or their features to the matcher (point 4), and retaining the one that maximizes the matching score $s_c(x)$, where x is the current image (or set of features) submitted by the attacker. In practice, it is a gradient-ascent technique that approximates the gradient of $s_c(x)$ numerically. In this case, the attacker is assumed to be able to observe $s_c(x)$ for any queried image, which may only be feasible if the system provides (or leaks) such information. Besides the aforementioned channel protection schemes, an additional defense mechanism consists of quantizing

[TABLE 1] THE CATEGORIZATION OF ATTACKS AND COUNTERMEASURES FOR BIOMETRIC SYSTEMS. FOR EACH ATTACK TECHNIQUE, WE ALSO REPORT THE TARGETED COMPONENT (ATTACK LOCATION) AND THE ATTACK POINT(S), ACCORDING TO FIGURE 1.

ATTACK TECHNIQUE	ATTACK LOCATION	ATTACK POINT(S)	DEFENSE
SPOOFING	SENSOR	1	LIVENESS DETECTION, MULTIBIOMETRICS (SECURE FUSION)
REPLAY	INTERFACES/CHANNELS	2, 4, 7	ENCRYPTED CHANNEL, TIME STAMP, CHALLENGE-RESPONSE, PHYSICAL ISOLATION
HILL CLIMBING	INTERFACES/CHANNELS	2, 4	ENCRYPTED CHANNEL, TIME STAMP, CHALLENGE-RESPONSE, PHYSICAL ISOLATION, SCORE QUANTIZATION
MALWARE INFECTION	MODULES/ALGORITHMS	3, 5, 8–11	SECURE CODE, SPECIALIZED HARDWARE, ALGORITHMIC INTEGRITY
TEMPLATE THEFT, SUBSTITUTION, AND DELETION	TEMPLATE DATABASE	6	TEMPLATE ENCRYPTION, CANCELABLE/REVOKABLE TEMPLATES

the matching score to provide less accurate information to the attacker. However, attacks based on more sophisticated black-box optimization techniques, suited to quantized objective functions, can also be considered to make these countermeasures ineffective [10].

The algorithmic implementations of the software modules (points 3, 5, and 8–11) may exhibit vulnerabilities that can be exploited by a skilled attacker through well-known hacking techniques (e.g., buffer overflow), to install malicious software, i.e., malware, including worms, trojan horses, etc. This threat can be avoided or mitigated by exploiting well-known programming practices, like secure code programming, or using specialized hardware to perform some critical operations [8]. A secure programming practice is to check algorithmic integrity, i.e., that each algorithm and function correctly handles any input parameter and never shows any unexpected behavior. For instance, if the matching algorithm expects a vector $x \in \mathbb{R}^d$ as input, and instead receives an input with a different format, is it going to crash or provide an output anyway? In the latter case, how is such an output handled by the subsequent modules? Does it lead to accepting by error the given claim as genuine or not?

Template theft, substitution, and deletion attacks target the template database (point 6). If templates are not protected properly, one may be able to steal them, and use them to create a spoof (i.e., a fake template), to perform a replay attack, or to impersonate the targeted client on a different system and perform other operations, e.g., searching on protected databases (function creep) [8]. Another possibility is to replace a template to impersonate a client without requiring any sophisticated attack as spoofing or replay; e.g., an attacker may add his own fingerprint template to the set of templates belonging to another client. Additionally, templates of a given client can be deleted to cause a denial of service, i.e., to avoid the targeted client to be recognized successfully. Countermeasures include template encryption, and also the use of cancelable/revokable templates, which can be used only on a specific system and reissued if stolen. The idea is to encode the templates using a key or pin code that can be changed to re-enroll the user and create a novel, different encrypted template [8].

ADVERSARIAL BIOMETRIC RECOGNITION

Here we analyze biometric system security in terms of our previously discussed framework by making assumptions on the adversary's goal, knowledge, capability, and attack strategy that are suited to biometric applications. Our aim is threefold: 1) to provide a well-structured categorization of the vulnerabilities of biometric systems and of the corresponding attacks, also through the definition of different, pertinent attack scenarios; 2) to provide a formal characterization of existing attacks within our framework and envision more sophisticated and effective attack strategies; and 3) to identify suitable countermeasures and defenses inspired by previous work on adversarial machine learning.

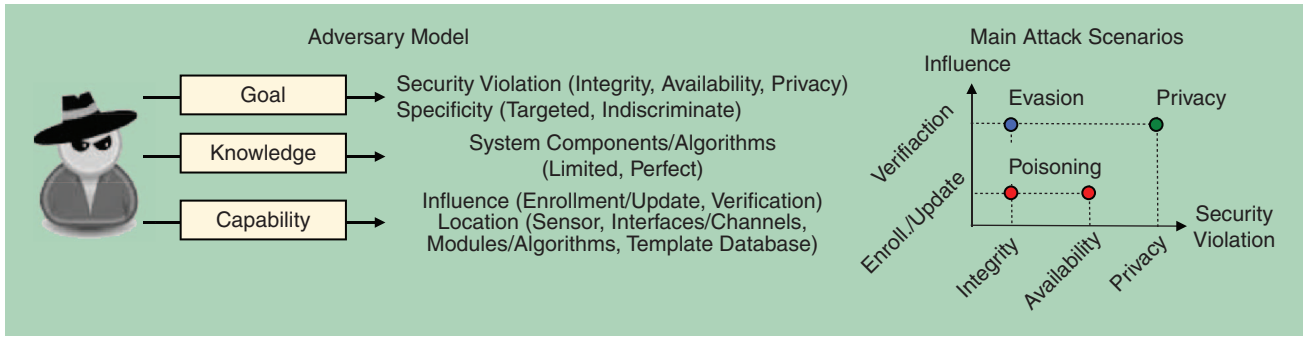
■ *Adversary's goal*: It is defined in terms of security violation and attack specificity. Biometric system security can be violated by an attacker that aims at impersonating a genuine user (integrity violation); at compromising the template galleries of genuine users to deny them access to the system, causing a

denial of service (availability violation); or at violating the privacy of genuine users, e.g., by inferring their templates through a hill-climbing attack (privacy violation). The attack specificity can be targeted, if the attack targets a specific set of clients, or indiscriminate, if any client may be affected.

■ *Adversary's knowledge*: It is defined by leveraging on the definition of the attack surface of biometric systems given in the previous section by making specific assumptions on what the attacker knows of the system components and how they work. According to Figure 1 and Table 1, the attacker may know 1) the kind of sensor used (point 1), e.g., an optical or capacitive fingerprint sensor; 2) which interfaces/channels are used to implement connections (points 2, 4, and 7), e.g., if an insecure channel over the Internet is used to send the acquired images to the feature extractor (point 2); 3) how the modules/algorithms work, and whether they are vulnerable or not (points 3, 5, and 8–11), in particular, the feature mapping ϕ (point 3), the matching algorithm s (point 5), the decision threshold t_c (point 8), the fusion rule s_c (point 9), and, if the template update is implemented, the self-update threshold θ_c (point 10) and the template update policy (point 11); and 4) some of the templates stored in the template database (point 6). The attacker may also be able to collect images of the same biometric traits using other techniques; e.g., acquiring latent fingerprints, or collecting face images of the targeted clients from social networks. From a machine-learning perspective, this amounts to having different levels of knowledge of the classifier's training data. In practice, it is worth noting that attackers typically have limited knowledge of the sensors and algorithms used, of the users' templates, and of any other system components (e.g., communication channels, template encryption schemes, etc.). Several previous works have, however, considered vulnerabilities of biometric systems without clearly pointing out the underlying assumptions on the adversary's knowledge required to perform the corresponding attack. Under our framework, such assumptions become clearly explicit.

■ *Adversary's capability*: This can also be defined in terms of the attack location: 1) the sensor (point 1); 2) interfaces/channels (points 2, 4, and 7); 3) the internals or even the output of modules and algorithms (points 3, 5, and 8–11), e.g., through malware infection attacks; and 4) the template database (point 6). In addition, one has to define the attack influence, i.e., the capability of manipulating the input data (e.g., using fake biometric traits), and how such data may be used to update the system (e.g., in adaptive biometric systems the attacker can produce spoofing attacks that can subsequently poison the clients' templates [20], [21]). Accordingly, the attack can influence only verification, or also enrollment/update.

■ *Attack strategy*: According to the adversary's goal [generally expressed in terms of an objective function $g(a; \theta)$], knowledge (given in terms of the parameter vector $\theta \in \Theta$) and capability (which defines the feasible set of attack strategies $a \in \mathcal{A}$), an optimal attack strategy can be defined to implement the attack, as explained before; see (1). For instance, assume that the attacker aims to impersonate an enrolled client (integrity targeted attack), and she is only able to acquire a latent fingerprint



[FIG2] A conceptual representation of the adversary model and of the main attack scenarios (given in terms of the corresponding security violation and attack specificity) according to our framework.

of the client, without having any other knowledge of the system components and algorithms. Then, the corresponding optimal attack strategy amounts to fabricating a fake fingerprint that is as similar as possible to the latent one, and using it to perform a spoofing attack. In this case, $g(a; \theta)$ can be regarded as a measure of the similarity between the fake and the latent fingerprint, as θ only contains information related to the latent fingerprint, and a corresponds to the fake fingerprint. A more skilled attacker may, however, also know the matching algorithm s and the fusion rule s_c used by the system, and may be able to collect more than a single fingerprint image of the targeted client. As an application example of our framework, we will show that, under this setting, more sophisticated and effective spoofing attacks can be fabricated. As another application example, we will also consider a poisoning attack against an adaptive face verification system and propose a novel countermeasure based on the sanitization of the client's templates.

In the following, we define a set of representative attack scenarios to categorize known attacks according to our framework. The framework and the considered attack scenarios are also represented in Figure 2.

CATEGORIZATION OF BIOMETRIC ATTACK SCENARIOS

Previous work has categorized attacks to biometric systems and countermeasures simply in terms of the attack points of Figure 1 (e.g., spoofing attacks to the sensor, countered by liveness detection

techniques, and attacks to a compromised channel, countered by channel encryption). Instead, looking at them from the broader perspective opened by our framework allows us to identify three main attack scenarios, described in the following, in which the attacks and countermeasures discussed in the previous sections play different roles. A few examples of known attack and defenses are categorized in Table 2 in terms of these attack scenarios. This also clears the way both to identify novel, more sophisticated attacks against biometric systems, and to adapt the corresponding countermeasures from the adversarial learning field.

EVASION

The goal of this attack scenario is to impersonate a client (integrity, targeted/indiscriminate attack). To this end, knowledge of the client's biometric trait is required, e.g., to create a fake trait or to carry out a replay attack. Attacks exploiting perfect knowledge of the targeted client's biometric trait include the so-called consensual method (in which the targeted client voluntarily provides the required biometric trait to the attacker), and template stealing. Conversely, exploiting a latent fingerprint is an example of limited knowledge, since the attacker may only partially know or observe the required biometric trait. A limited knowledge about the rest of the biometric system can also be sufficient. In this scenario, the capability of the attacker consists of manipulating data during the verification step, whereas no influence on the enrollment/update step is assumed (in particular, she can not access the template database). Most frequently, the attack

[TABLE 2] EXAMPLES OF CATEGORIZATION OF PREVIOUS WORK ON BIOMETRIC SECURITY ACCORDING TO THE THREE MAIN ATTACK SCENARIOS DEFINED IN ADVERSARIAL MACHINE LEARNING: EVASION, POISONING, AND PRIVACY ATTACKS.

	GOAL		KNOWLEDGE	CAPABILITY		ATTACK STRATEGY
	VIOLATION	SPECIFICITY		INFLUENCE	LOCATION	
EVASION ATTACKS						
MATSUMOTO et al. [23], RODRIGUES et al. [9], JOHNSON et al. [24]	INTEGRITY	TARGETED, INDISCRIMINATE	PERFECT (CONSENSUAL FAKE)	VERIFICATION	SENSOR	SPOOFING
POISONING ATTACKS						
BIGGIO et al. [21, 20]	INTEGRITY, AVAILABILITY	TARGETED, INDISCRIMINATE	PERFECT, LIMITED (UNKNOWN TEMPLATES)	ENROLLMENT/ UPDATE	SENSOR	SPOOFING (FACE)
PRIVACY ATTACKS						
ADLER [10], GALBALLY et al. [11], MARTINEZ et al. [12]	PRIVACY	TARGETED, INDISCRIMINATE	LIMITED (UNKNOWN TEMPLATES)	VERIFICATION	MATCHER	HILL CLIMBING

strategy corresponding to an evasion attack consists of submitting a fake trait (spoof) to the sensor (point 1), or of replaying the acquired image (point 2) into the system. In rarer cases (disregarded here), the biometric system can be infected by malware, potentially allowing the attacker to arbitrarily manipulate the functionality or the output of any system component.

POISONING

This nontrivial attack scenario has been originally defined in the context of adaptive biometric systems in our previous work [20], [21], inspired by our adversarial learning framework, and by work on poisoning learning algorithms [1]–[4]. The goal of poisoning attacks can be either an integrity or availability security violation; it can be either targeted to a specific client, or indiscriminate (see the section “Poisoning Biometric Systems that Learn from Examples”). The adversary’s knowledge can be perfect or limited, depending on whether each of the system’s components is exactly known to the attacker. More precisely, the attacker may have perfect (or limited) knowledge of each of the components discussed in Figure 1, including the targeted clients’ templates, the matching algorithm, the template update algorithm, and the decision and update thresholds. The attacker’s capability consists of modifying the template database, either by directly manipulating it (e.g., through malware infection), or, more realistically, by submitting fake traits that are erroneously used to update the template gallery of a given client. In terms of security violation, an integrity violation thus amounts to replacing a victim’s template with an attacker’s template or to adding an attacker’s template in the victim’s gallery. This indeed allows the attacker to impersonate the victim without using any further spoofing or replay attack, but directly using her own biometric trait. The goal of an availability violation is to cause a denial of service, instead, by replacing or compromising the majority of templates in the victim’s gallery. This will indeed deny the victim access to the system. Under this setting, the attack strategy amounts to compromising the template gallery either by introducing an attacker’s template in the victim’s gallery (i.e., integrity violation), or by compromising the maximum number of victim’s templates (i.e., availability violation). If the template database can not be compromised directly, the attacker can produce a well-crafted sequence of fake traits to gradually drift the victim’s template gallery toward the desired set of templates, while minimizing the number of fake traits required to complete the attack. An example of such an attack is given next.

PRIVACY

In this case, the goal is to retrieve confidential information (i.e., one or more templates) about either a given set of clients (targeted attack) or about any client (indiscriminate attack). This is typically a preliminary step before performing another kind of attack (evasion or poisoning), when no simpler way to retrieve information on the victims’ templates exists (e.g., acquiring a face image through a social network, or a latent fingerprint). To this end, the attacker can gain knowledge from the system’s feedback, e.g., the outcome of the verification decision (either accept or reject), or the score value $s_c(x)$ (as in hill-climbing attacks). In common settings (i.e., disregarding cases like malware infection), the capability consists of

sending a number of query images through a remote channel and observing the available feedback; e.g., if the sensor and the matcher are remotely operating, and interconnected through the Internet, an attacker may perform a man-in-the-middle attack and send replayed images through the channel. It is clear that the attack strategy in this case corresponds to an hill-climbing attack.

SECURE-BY-DESIGN BIOMETRIC SYSTEMS

The considered adversary model can be exploited not only to provide a different categorization of known defense mechanisms for biometric recognition systems but also to identify novel countermeasures among those proposed in adversarial learning, which can help countering attacks against machine-learning and pattern recognition algorithms used in biometric systems. We discuss them below, with reference to the three aforementioned attack scenarios.

COUNTERING EVASION

In this scenario, the main attack strategies involve spoofing and replay. As reported in Table 1, the pertinent defenses are: liveness detection, multibiometric systems with secure score-level fusion rules, encrypted channels and time-stamp/challenge-response schemes, and physical isolation (e.g., match-on-smart-cards). Novel defense mechanisms can also be devised, inspired by the adversarial machine-learning field. In particular, to counter evasion attacks, one can consider secure learning techniques. They consist of modifying existing learning algorithms (and developing novel ones) that explicitly take into account a specific kind of adversarial data manipulation. They follow the paradigm of security by design, which advocates that a system should be designed from the ground up to be secure. In the context of biometric systems, secure learning techniques can be exploited to design trainable score-level fusion rules, such as those based on game theory, or on the framework of learning with invariances [14], [25]–[27]. Investigating this issue would be an interesting research direction for future work.

COUNTERING POISONING

Spoofing and replay are the main attack strategies that are also under this scenario and, thus, the same defenses listed in Table 1 can be also exploited in this case. In addition, other countermeasures can be considered, among those proposed in adversarial learning, to improve the security of the training phase in the presence of poisoning, which may occur when the system is retrained on data collected during operation [15], [28]. These include secure learning (similar to countermeasures to evasion) and data sanitization. In a biometric system, the latter consists of detecting outlying template updates that may compromise the template gallery of a given client, e.g., by adding an impostor’s template to the targeted client’s gallery, or by replacing some of the client’s templates. We will give a concrete example of a novel defense based on template sanitization in the next section. We point out that these additional defenses can be considered complementary to those listed in Table 1, like liveness detection and channel encryption.

PRESERVING PRIVACY

Known defenses that can be exploited against attacks targeting the template database are mainly based on template encryption

schemes [8] (Table 1). Score quantization has been also proposed to counter hill-climbing attacks, but it has already been shown to be ineffective [10]. Moreover, attacks proposed in adversarial learning have already been capable of reverse engineering the classifier by only exploiting only feedback on its decisions [16], [17]; thus, even by only looking at genuine or impostor classifications, an attacker may be able to successfully perform a hill-climbing attack. Among the proposed countermeasures that have not yet been considered for biometric systems, it would be worth investigating in future work the ones based on randomization and disinformation. They follow the paradigm of security by obscurity, aiming to improve system security by hiding information to the attacker. They have been suggested in adversarial learning to counter reverse-engineering attacks. This can be achieved by denying access to the actual classifier or training data, and randomizing the classifier's output to give imperfect feedback to the attacker [1]–[4], [29].

APPLICATION EXAMPLES

Here we consider two application examples of our framework related to the development of sophisticated spoofing and poisoning attacks against face-verification systems, respectively.

IMPROVED FACE SPOOFING FROM MULTIPLE FACES

Let us assume we are given a face verification system that authenticates clients by matching the acquired face image against the template gallery of the claimed identity (consisting of n images acquired during enrollment), and then thresholding the corresponding average score. According to the architecture depicted in Figure 1, we assume that our system maps the submitted face image $z \in \mathcal{Z}$ onto a reduced vector space \mathcal{X} using principal component analysis (PCA), and computes the matching score for client c as $s_c(x) = (1/n) \sum_{i=1}^n s(x, x_i)$, where $s(x, x_i) = \exp\{-\|x - x_i\|\}$, and x_i is the i th template of the claimed identity. We further assume the attack scenario detailed below.

- *Adversary's goal*: The attacker aims to impersonate a targeted client (integrity, targeted attack).
- *Adversary's knowledge*: She is assumed to know 1) the feature extraction algorithm, 2) the matching algorithm s , 3) the fusion rule s_c , and 4) a set of n face images $\{\hat{x}_j\}_{j=1}^n$ of the targeted client, different from those in the client's template gallery (e.g., potentially collected from a social network).
- *Adversary's capability*: She can only submit printed photos of faces to the sensor during verification.

■ *Adversary's strategy*: Under these assumptions, the attacker can approximate the score $s_c(x)$ computed by the targeted system for the claimed identity c using the collected face images of the victim, i.e., she can compute an estimate $\hat{s}_c(x) = (1/n) \sum_{j=1}^n s(x, \hat{x}_j)$. Accordingly, the optimal attack strategy is given by:

$$x^* = \underset{x}{\operatorname{argmax}} \hat{s}_c(x) = \frac{1}{n} \sum_{j=1}^n s(x, \hat{x}_j), \quad (2)$$

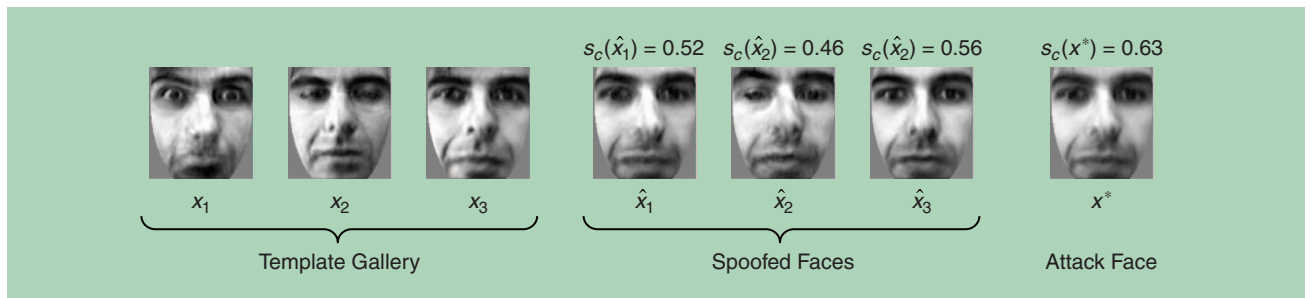
where \hat{s}_c is the attacker's goal function $g(a, \theta)$ (1), and x^* is the attack sample that maximizes the objective in the PCA-induced feature space. The above problem can be solved by a simple gradient-ascent algorithm, and the resulting attack sample x^* can then be projected back onto the space of face images \mathcal{Z} (where each feature corresponds to the gray-level value of a pixel) by inverting the PCA-induced mapping. This is also possible if more sophisticated matching algorithms and feature representations are used, using well-crafted heuristics. Please refer to [21] for further details.

An example of this improved spoofing technique on a simple case involving $n = 3$ templates is shown in Figure 3, where we also report the values of s_c for each of the client's face images $\{\hat{x}_j\}_{j=1}^n$. It can be seen that the final attack face x^* yields a higher probability (i.e., s_c value) of successfully impersonating the victim than any of the available images $\{\hat{x}_j\}_{j=1}^n$.

POISONING BIOMETRIC SYSTEMS THAT LEARN FROM EXAMPLES

We now report a different application example focusing on a poisoning attack against an adaptive face recognition system, and on the development of a corresponding defense. In recent work [20], [21], we have shown that an attacker may exploit the system's adaptation mechanism to compromise the templates of a given client by presenting a well-crafted sequence of fake faces to the camera, with the goal of denying him access to the system. At the same time, if the attacker replaces the targeted client's templates with her templates, she may also impersonate the client without presenting any fake traits to the sensor.

The face verification system considered in this example, as in [20] and [21], authenticates clients based on matching the acquired face image with a stored average template, referred to as *centroid*. For each client, the centroid is updated using the self-update algorithm: if the submitted face image is similar enough to the centroid, the

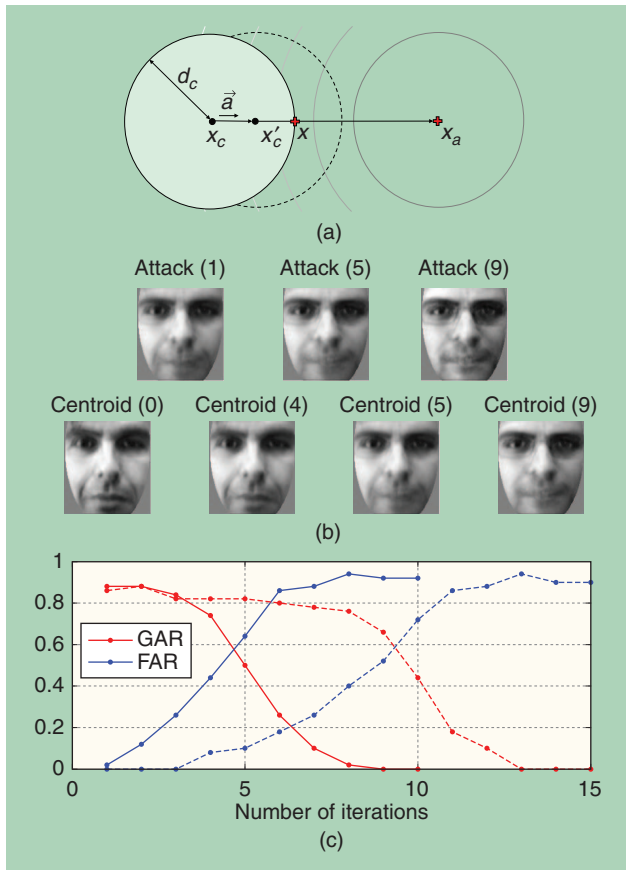


[FIG3] Face spoofing from multiple images. The client's templates $\{x_i\}_{i=1}^3$, the spoofed faces $\{\hat{x}_j\}_{j=1}^3$, and the final attack face x^* [obtained from solving (2)] are shown, along with the corresponding s_c values.

latter is updated incorporating the new image into the computation of the average face image. As in the previous case, we consider a PCA-based mapping to map face images from \mathcal{Z} onto a reduced vector space \mathcal{X} . The matching score for client c is computed here as $s_c(x) = \exp\{-\|x - x_c\|^2\}$, where x_c is the client's centroid. The centroid x_c is initially computed as the average of n templates and, when $s_c(x) \geq \theta_c$, updated as $x'_c = x_c + (1/n)(x - x_c)$, i.e., slightly drifted toward x . Accordingly, in the PCA-based feature space, x_c is updated if the acquired image x is within a hypersphere centered on x_c , with radius d_c dependent on the update threshold θ_c . The complete attack scenario is given next.

■ **Adversary's goal:** It is that of replacing the centroid x_c of a given client with an attacker's template x_a , both to deny access to client c , and to allow the attacker to impersonate c using her own face (i.e., a targeted attack violating both system availability and integrity).

■ **Adversary's knowledge:** The attacker is assumed to know 1) the feature extraction algorithm; 2) the matching algorithm; 3) the template update algorithm; and 4) the decision and self-update threshold. In the case of perfect knowledge,

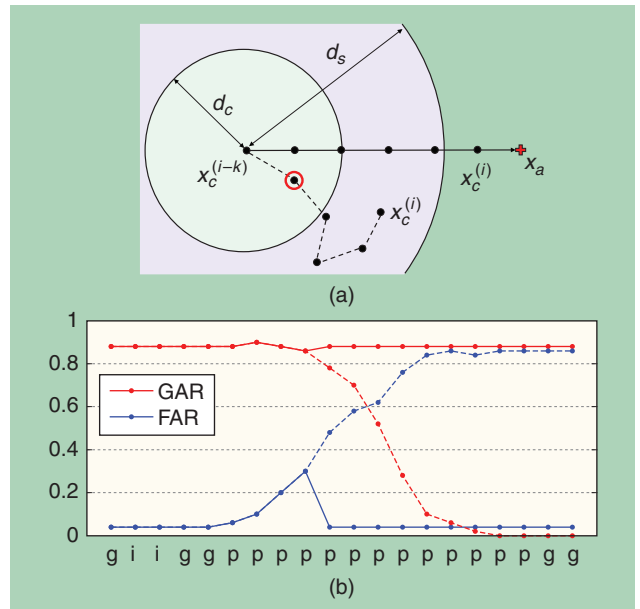


[FIG4] (a) A poisoning attack with perfect knowledge. The circles centered on x_a represent the objective function $\|x - x_a\|$, minimized by the attack point x on the feasible domain $\|x - x_c\| \leq d_c$. The updated centroid x'_c and the feasible domain for the next attack iteration are also shown. (b) Attack samples and the victim's centroids for poisoning with perfect knowledge at different iterations. (c) Genuine acceptance rate (GAR) and false acceptance rate (FAR) values in the presence (solid lines) and in the absence (dashed lines) of template sanitization, after different centroid updates, including genuine ("g") and impostor ("i") attempts, and poisoning attacks ("p") with perfect knowledge.

she also knows the centroid x_c of the targeted client c , while when limited knowledge is considered, only a good estimate of x_c is available to the attacker, e.g., a frontal face image of the victim collected from a social network.

■ **Adversary's capability:** The attacker can modify the template database by presenting fake faces at the sensor that enable template self-update. The attack influence is thus over the enrollment/update phase.

■ **Attack strategy:** Under these assumptions, the shortest sequence of fake traits required to replace the victim's centroid x_c with that of the attacker x_a can be found by solving the following optimization problem, for each sample x in the attack sequence: $\min_x \|x - x_a\|^2$, subject to the update condition $\|x - x_c\| \leq d_c$ [see Figure 4(a)]. At each iteration, this amounts to finding the closest attack sample to x_a that enables an update of x_c . The solution is simply given as $x = x_c + d_c \bar{a}$, where $\bar{a} = (x_a - x_c / \|x_a - x_c\|)$ is the so-called attack direction. In practice, each attack sample x is found at the intersection between the hypersphere corresponding to the update condition, and the line connecting x_a and x_c . As in the previous example, the face images for the attack sequence can be obtained by projecting the attack samples from the PCA-induced space \mathcal{X} onto the space of face images \mathcal{Z} . Then, the attacker can fabricate the corresponding fake faces (e.g., by printing them on paper), and present them in the right order to the sensor. An example of how the victim's centroid is gradually updated by the corresponding sequence of attack faces, under the assumption of perfect



[FIG5] (a) Template sanitization. If the current centroid $x_c^{(i)}$ falls outside the sanitization hypersphere (dark gray area), as for the poisoning attack sequence (solid line), the centroid $x_c^{(i-k)}$ is restored; otherwise, as for the hypothesized genuine update sequence (dashed line), the center of the sanitization hypersphere is updated to $x_c^{(i-k+1)}$ (red circled point). (b) GAR and FAR values in the presence (solid lines) and in the absence (dashed lines) of template sanitization, after different centroid updates, including genuine ("g") and impostor ("i") attempts, and poisoning attacks ("p") with perfect knowledge.

knowledge, and for $n = 5$, is given in Figure 4(b). In (c), we show how the GAR and FAR vary as the attack proceeds, under perfect and limited knowledge of the victim's template. Note how the probability of authenticating the attacker (without presenting any fake face) as the victim (i.e., the FAR) increases, while the probability of correctly authenticating the victim as a genuine user (i.e., the GAR) decreases, since the victim's template is gradually morphed toward the attacker's face while the attack is in progress. Results for the perfect and limited knowledge attacks are similar, despite the latter case requires more iterations (i.e., submitting more fake faces) to compensate the lack of knowledge of the victim's template. The exact number of iterations required to complete both attacks can also be analytically computed [21].

■ *Template sanitization*: Here we present a novel countermeasure based on the idea of sanitizing the template gallery (i.e., identifying anomalous template updates), inspired by the countermeasures proposed in adversarial machine learning against poisoning attacks [1]–[4], [28].

The underlying idea is to analyze whether the sequence of the most recent k updated centroids $\mathbf{x}_c^{(i-k)}, \dots, \mathbf{x}_c^{(i)}$ (where i denotes the current iteration) falls within a given region of the feature space, called sanitization hypersphere [see Figure 5(a)]. If the current centroid $\mathbf{x}_c^{(i)}$ falls within the sanitization hypersphere, i.e., if $\|\mathbf{x}_c^{(i)} - \mathbf{x}_c^{(i-k)}\| \leq d_s$, then the center of the sanitization hypersphere is updated to the next centroid in the sequence, i.e., $\mathbf{x}_c^{(i-k+1)}$; otherwise, the current center of the sanitization hypersphere $\mathbf{x}_c^{(i-k)}$ is restored as the current centroid. In this case, an alert may be also (or alternatively) raised to the system administrator to report the anomalous update. The rationale of this approach is to identify sequences of centroid updates that consistently drift the centroid toward a given, biased direction, within a small number of iterations (as it happens in the presence of a poisoning attack), assuming that genuine updates exhibit a different (i.e., less biased and more random) behavior. The parameter k and the hypersphere radius d_s of the proposed approach should thus be chosen such that $d_s \leq k(d_c/n)$, otherwise poisoning attacks will not be detected, as they drift the centroid of an amount equal to (d_c/n) at each iteration.

In Figure 5(b), we report an example using the same attacker and victim pair considered in the previous case. Initially, we simulate a number of random accesses to the system, including genuine and impostor attempts, which do not significantly affect GAR and FAR. Then, the attacker launches a poisoning attack, and, as in the previous case, the GAR decreases and the FAR increases. If template sanitization is not implemented, the attack succeeds, and system integrity and availability are compromised. Conversely, in the presence of template sanitization, the attack is detected after four iterations, and a previous centroid is restored. This avoids the attack to succeed, preserving normal system operation and security. Although this simple countermeasure can be misled by a poisoning attack in which the attack samples are closer to the current centroid (instead of lying exactly at the boundary of the feasible domain), this would require the attacker to perform a significantly higher number of iterations to complete

the attack. We can thus conclude that the proposed sanitization technique helps improving system security.

SUMMARY AND OPEN PROBLEMS

In the last decade, the use of electronic devices in our daily lives has become increasingly pervasive, providing several advantages in managing our tasks and communicating with other people. However, with such a huge number of powerful devices connected to the so-called Internet of Things (e.g., smartphones, smart TVs, etc.), the number of potential attack points and vulnerabilities has significantly increased, as well as chances for attackers to compromise the corresponding devices and systems. In addition, the level of sophistication of attacks has become increasingly higher over the years, witnessing the presence of very skilled attackers and strong economic incentives behind these activities. Biometrics can be considered a potential tool for improving the security of such systems in the digital era. However, besides having a strong deterrent effect, they should really be designed to be intrinsically secure, to successfully resist the very sophisticated attacks that may be incurred during operation.

In this article, we have provided an overview of the current state of the art on biometric system security from a perspective inspired by the field of adversarial machine learning. We have discussed how this novel perspective may not only inspire the simulation of more sophisticated attack scenarios, but also how, based on such scenarios, more effective countermeasures can be proactively developed. As concrete application examples, we have considered a sophisticated spoofing attack and a poisoning attack against an adaptive face verification system in which the attacker gradually compromises the template gallery of a given client by presenting a well-crafted sequence of fake faces. We have also proposed a novel countermeasure based on template sanitization. Another example of how the proposed perspective based on adversarial machine learning may depict novel attack scenarios and inspire potential countermeasures is related to recent work in adversarial learning, which has shown that clustering algorithms may be significantly vulnerable to well-planned attacks [30]. In the biometric setting, similar attacks may target systems that perform template selection and update exploiting clustering algorithms. Although investigating and developing secure clustering algorithms against adaptive and intelligent attackers is still an open issue in the adversarial machine-learning field, the corresponding results may also inspire countermeasures that can be adapted to improve the security of template selection and update procedures in biometric systems.

ACKNOWLEDGMENTS

This work has been partly supported by the TABULA RASA project, 7th Framework Research Programme of the European Union, grant agreement number 257289; and by the project CRP-18293, funded by Regione Autonoma della Sardegna, L.R. 7/2007, Bando 2009.

AUTHORS

Battista Biggio (battista.biggio@diee.unica.it) received the M.S. degree in electronic engineering (with honors) and the Ph.D. degree in electronic engineering and computer science, respectively, in 2006 and 2010, from the University of Cagliari, Italy,

where he has been with the Department of Electrical and Electronic Engineering of the same university since 2007 and now holds a postdoctoral position. His research interests currently include adversarial learning, multiple classifier systems, kernel methods, biometric authentication, spam, and malware detection. He serves as a reviewer for the main international conferences and journals in these fields. He is a member of the IEEE.

Giorgio Fumera (fumera@diee.unica.it) received the M.S. degree in electronic engineering (with honors) and the Ph.D. degree in electronic engineering and computer science, respectively in 1997 and 2002, from the University of Cagliari, Italy. Since February 2010, he has been an associate professor of computer engineering in the Department of Electrical and Electronic Engineering at the same university. His research interests are related to methodologies and applications of statistical pattern recognition, and include multiple classifier systems, classification with the reject option, adversarial classification, and document categorization. He acts as reviewer for the main international conferences and journals in these fields. He is a Member of the IEEE.

Paolo Russu (paolo.russu@diee.unica.it) received the B.S. degree and M.S. degree (with honors) in 2010 and 2013, respectively, from the University of Cagliari, Italy. Since May 2014 he has been a Ph.D. student at the same university. His research focuses on adversarial machine learning, biometrics, and computer security.

Luca Didaci (luca.didaci@diee.unica.it) received the M.S. degree in electronic engineering and the Ph.D. degree in electronic engineering and computer science, respectively, in 2001 and 2005, from the University of Cagliari, Italy. Since 2002 he has been an assistant professor of computer engineering in the Department of Electrical and Electronic Engineering at the same university. His research interests are in the fields of pattern recognition and applications. In particular, they rely on semisupervised multiple classifier systems, semisupervised biometric systems, and adversarial classification. He acts as reviewer for the main international conferences and journals in these fields. He is a Member of the IEEE.

Fabio Roli (roli@diee.unica.it) received his M.S. degree (with honors) and Ph.D. degree in electronic engineering from the University of Genoa, Italy. He was a member of the research group on Image Processing and Understanding of the University of Genoa, Italy, from 1988 to 1994. He was adjunct professor at the University of Trento, Italy, in 1993 and 1994. In 1995, he joined the Department of Electrical and Electronic Engineering of the University of Cagliari, Italy, where he is now professor of computer engineering and head of the research group on pattern recognition and applications. His research over the past 20 years addressed the design of pattern recognition systems in real applications. He is Fellow of the IEEE and of the International Association for Pattern Recognition.

REFERENCES

[1] M. Barreno, B. Nelson, R. Sears, A. D. Joseph, and J. D. Tygar, "Can machine learning be secure?" in *Proc. ACM Symp. Information, Computer and Communications Security (ASIACCS'06)*. New York: ACM, 2006, pp. 16–25.

[2] L. Huang, A. D. Joseph, B. Nelson, B. Rubinstein, and J. D. Tygar, "Adversarial machine learning," in *Proc. 4th ACM Workshop Artificial Intelligence and Security* Chicago, IL, 2011, pp. 43–57.

[3] B. Biggio, G. Fumera, and F. Roli, "Security evaluation of pattern classifiers under attack," *IEEE Trans. Knowledge Data Eng.*, vol. 26, no. 4, pp. 984–996, 2014.

[4] B. Biggio, G. Fumera, and F. Roli, "Pattern recognition systems under attack: Design issues and research challenges," *Int. J. Pattern Recogn. Artif. Intell.*, vol. 28, no. 7, p. 1460002, 2014.

[5] P. Laskov and R. Lippmann, Eds., *NIPS Workshop on Machine Learning in Advances in Environments for Computer Security*, 2007.

[6] A. D. Joseph, P. Laskov, F. Roli, and D. Tygar, Eds., *Dagstuhl Perspectives Workshop on Machine Learning Methods for Computer Security*, 2012.

[7] N. K. Ratha, J. H. Connell, and R. M. Bolle, "An analysis of minutiae matching strength," in *AVBPA*, ser. LNCS, J. Bigün and F. Smeraldi, Eds., vol. 2091. New York: Springer, 2001, pp. 223–228.

[8] A. K. Jain, K. Nandakumar, and A. Nagar, "Biometric template security," *J. Adv. Signal Process.*, vol. 2008, New York: Hindawi, pp. 1–17, 2008.

[9] R. N. Rodrigues, L. L. Ling, and V. Govindaraju, "Robustness of multimodal biometric fusion methods against spoof attacks," *J. Vis. Lang. Comput.*, vol. 20, no. 3, pp. 169–179, 2009.

[10] A. Adler, "Vulnerabilities in biometric encryption systems," in *Proc. 5th Int. Conf. Audio- and Video-Based Biometric Person Authentication*, ser. LNCS, T. Kanad, A. Jain, N. K. Ratha, Eds., vol. 3546. Rye, NY: Springer, 2005, pp. 1100–1109.

[11] J. Galbally, C. McCool, J. Fierrez, S. Marcel, and J. Ortega-Garcia, "On the vulnerability of face verification systems to hill-climbing attacks," *Pattern Recogn.*, vol. 43, no. 3, pp. 1027–1038, 2010.

[12] M. Martinez-Diaz, J. Fierrez, J. Galbally, and J. Ortega-Garcia, "An evaluation of indirect attacks and countermeasures in fingerprint verification systems," *Pattern Rec. Lett.*, vol. 32, no. 12, pp. 1643–1651, 2011.

[13] B. Biggio, I. Corona, D. Maiorca, B. Nelson, N. Šrnđić, P. Laskov, G. Giacinto, and F. Roli, "Evasion attacks against machine learning at test time," in *European Conf. Machine Learning and Principles and Practice of Knowledge Discovery in Databases, Part III*, ser. LNCS, H. Blockeel, K. Kersting, S. Nijssen, and F. Železný Eds., vol. 8190. Berlin: Springer, 2013, pp. 387–402.

[14] A. Globerson and S. T. Roweis, "Nightmare at test time: Robust learning by feature deletion," in *Proc. 23rd Int. Conf. Machine Learning*, W. W. Cohen and A. Moore, Eds., vol. 148. New York: ACM, 2006, pp. 353–360.

[15] B. I. Rubinstein, B. Nelson, L. Huang, A. D. Joseph, S.-h. Lau, S. Rao, N. Taft, and J. D. Tygar, "Antidote: Understanding and defending against poisoning of anomaly detectors," in *Proc. 9th ACM SIGCOMM Internet Measurement Conf., IMC'09*. New York: ACM, 2009, pp. 1–14.

[16] D. Lowd and C. Meek, "Adversarial learning," in *Proc. 11th ACM SIGKDD Int. Conf. Knowledge Discovery Data Mining*. Chicago, IL: ACM, 2005, pp. 641–647.

[17] B. Nelson, B. I. Rubinstein, L. Huang, A. D. Joseph, S. J. Lee, S. Rao, and J. D. Tygar, "Query strategies for evading convex-inducing classifiers," *J. Mach. Learn. Res.*, vol. 13, pp. 1293–1332, May 2012.

[18] U. Uludag, A. Ross, and A. K. Jain, "Biometric template selection and update: A case study in fingerprints," *Pattern Recogn.*, vol. 37, no. 7, pp. 1533–1542, 2004.

[19] C. Ryu, H. Kim, and A. K. Jain, "Template adaptation based fingerprint verification," in *Proc. 18th Int. Conf. Pattern Recognition, ICPR'06*, vol. 4, Washington, DC, 2006, pp. 582–585.

[20] B. Biggio, G. Fumera, F. Roli, and L. Didaci, "Poisoning adaptive biometric systems," in *Structural, Syntactic, and Statistical Pattern Recognition*, LNCS, G. Gimel'farb, E. Hancock, A. Imiya, A. Kuijper, M. Kudo, S. Omachi, T. Windeatt, and K. Yamada, Eds. New York: Springer, 2012, vol. 7626, pp. 417–425.

[21] B. Biggio, L. Didaci, G. Fumera, and F. Roli, "Poisoning attacks to compromise face templates," in *Proc. 6th IAPR Int. Conf. Biometrics*, Madrid, Spain, 2013, pp. 1–7.

[22] B. Biggio, Z. Akhtar, G. Fumera, G. L. Marcialis, and F. Roli, "Security evaluation of biometric authentication systems under real spoofing attacks," *IET Biomet.*, vol. 1, no. 1, pp. 11–24, 2012.

[23] T. Matsumoto, H. Matsumoto, K. Yamada, and S. Hoshino, "Impact of artificial 'gummy' fingers on fingerprint systems," *Datenschutz und Datensicherheit*, vol. 26, no. 8, pp. 1–15, 2002.

[24] P. Johnson, B. Tan, and S. Schuckers, "Multimodal fusion vulnerability to non-zero effort (spoof) imposters," in *Proc. IEEE Int. Workshop on Information Forensics and Security*, 2010, pp. 1–5.

[25] N. Dalvi, P. Domingos, Mausam, S. Sanjay, and D. Verma, "Adversarial classification," in *Proc. 10th ACM SIGKDD Int. Conf. Knowledge Discovery Data Mining*, Seattle, 2004, pp. 99–108.

[26] M. Brückner, C. Kanzow, and T. Scheffer, "Static prediction games for adversarial learning problems," *J. Mach. Learn. Res.*, vol. 13, pp. 2617–2654, Sept. 2012.

[27] C. H. Teo, A. Globerson, S. Roweis, and A. Smola, "Convex learning with invariances," in *Proc. Neural Information Processing Systems 20*, J. C. Platt, D. Koller, Y. Singer, and S. T. Roweis, Eds. Cambridge, MA: MIT Press, 2008, pp. 1489–1496.

[28] G. F. Cretu, A. Stavrou, M. E. Locasto, S. J. Stolfo, and A. D. Keromytis, "Casting out demons: Sanitizing training data for anomaly sensors," in *Proc. IEEE Symp. Security and Privacy*. Los Alamitos, CA: 2008, pp. 81–95.

[29] A. D. Joseph, P. Laskov, F. Roli, J. D. Tygar, and B. Nelson, "Machine learning methods for computer security (Dagstuhl Perspectives Workshop 12371)," *Dagstuhl Manifestos*, vol. 3, no. 1, pp. 1–30, 2013.

[30] B. Biggio, I. Pillai, S. R. Bulò, D. Ariu, M. Pelillo, and F. Roli, "Is data clustering in adversarial settings secure?" in *Proc. ACM Workshop on Artificial Intelligence Security (AISeC'13)*. New York: ACM, 2013, pp. 87–98.