Preprint of the paper available at https://doi.org/10.1016/j.patrec.2020.10.008

Pattern Recognition Letters
journal homepage: www.elsevier.com

Cross-Resolution Face Recognition Adversarial Attacks

Fabio Valerio Massoli[a,**], Fabrizio Falchi[a], Giuseppe Amato[a]

[a]ISTI-CNR, via G. Moruzzi 1, 56124 Pisa, Italy

ABSTRACT

Face Recognition is among the best examples of computer vision problems where the supremacy of deep learning techniques compared to standard ones is undeniable. Unfortunately, it has been shown that they are vulnerable to adversarial examples - input images to which a human imperceptible perturbation is added to lead a learning model to output a wrong prediction. Moreover, in applications such as biometric systems and forensics, cross-resolution scenarios are easily met with a non-negligible impact on the recognition performance and adversary's success. Despite the existence of such vulnerabilities set a harsh limit to the spread of deep learning-based face recognition systems to real–world applications, a comprehensive analysis of their behavior when threatened in a cross-resolution setting is missing in the literature. In this context, we posit our study, where we harness several of the strongest adversarial attacks against deep learning-based face recognition systems considering the cross-resolution domain. To craft adversarial instances, we exploit attacks based on three different metrics, i.e., L_1, L_2, and L_∞, and we study the resilience of the models across resolutions. We then evaluate the performance of the systems against the face identification protocol, open- and close-set. In our study, we find that the deep representation attacks represents a much dangerous menace to a face recognition system than the ones based on the classification output independently from the used metric. Furthermore, we notice that the input image's resolution has a non-negligible impact on an adversary's success in deceiving a learning model. Finally, by comparing the performance of the threatened networks under analysis, we show how they can benefit from a cross-resolution training approach in terms of resilience to adversarial attacks.

1. Introduction

Face Recognition (Wang and Deng, 2018; Deng et al., 2019) (FR) represents one of the most astonishing applications of Neural Networks (NNs), especially considering Deep Convolutional Neural Networks (DCNNs), that ultimately overcame standard computer vision techniques such as Gabor-Fisher (Liu and Wechsler, 2002) and local binary patterns (Ahonen et al., 2006). The study of such a problem began in the early 90s when Turk and Pentland (1991) proposed the Eigenfaces approach, and it only required two decades for Deep Learning (DL) approaches to start to dominate the field reaching recognition performance up to 99.80% Wang and Deng (2018), thus overcoming human ability. DL-based FR systems do not exploit the output of a classifier directly. Instead, they leverage the representation power (LeCun et al., 2015) of the learning models to extract face descriptors, i.e., multidimensional vectors, also called deep features or deep representations, to fulfill the recognition task.

Although FR systems obtain very high performance when trained with datasets comprising images acquired under controlled conditions, e.g., high-resolution, they suffer a drastic drop in reliability when tested against cross-resolution (CR) scenarios (Massoli et al., 2019) that naturally arise, for example, in surveillance applications (Zou and Yuen, 2011; Amato et al., 2019; Cheng et al., 2018). To counteract such a weakness, Ekenel and Sankur (2005) and Luo et al. (2019) proposed approaches that were not based on NNs. Instead, only recently such a problem has been tackled in the DL field (Massoli et al., 2020; Zhang et al., 2018).

To make the situation even worse, recently Szegedy et al. (2013); Biggio et al. (2013) showed that DL models are vulner-

**Corresponding author
e-mail: fabio.massoli@isti.cnr.it (Fabio Valerio Massoli)

able to the so-called *adversarial* examples - images to which a specific amount of noise, undetectable to humans, is added to induce a NN to output a wrong prediction. Unfortunately, the ability of an insightful adversary to jeopardize these learning models, considering both the digital (Dong et al., 2019; Song et al., 2018; Qiu et al., 2019; Kakizaki and Yoshida, 2019; Goswami et al., 2018) and physical (Sharif et al., 2016; Kurakin et al., 2016) domains, represents a significant concern in security-related applications such as DL-based biometrics systems (Sundararajan and Woodard, 2018) and forensics (Spaun, 2011). Thus, limiting their adoption in these fields.

In this context, we posit our contribution that we summarize as follows: i) we threaten two DCNNs by exploiting adversarial attacks based on three different metrics, i.e., L_1, L_2, and L_∞; ii) we generate attacks not only towards a classification objective but also against a similarity one. Indeed, FR systems typically do not exploit a DCNN classification output. Instead, they leverage the ability of NNs to generate discriminative deep representations among which a similarity criterion is evaluated to fulfill the recognition task; iii) we conduct the attacks in a cross-resolution domain, thus emulating a real-world scenario for an FR system; iv) we analyze the success rates of the various attacks across resolutions, studying if a DL model can benefit from a cross-resolution training procedure in terms of robustness to adversarial attacks; v) we analyze the robustness of the models through the face identification protocol (Grother et al., 2019) considering both the open- and close-set settings.

The rest of the paper is structured as follows. In Section 2, we briefly present some related works, while in Section 3, we describe the attacks algorithms we use. Subsequently, in section 4, we explain our experimental procedure and the dataset we use, while in Section 5, we present the results from the experimental campaign. Finally, in Section 6, we report our conclusions.

2. Related Works

To the best of our knowledge, this is the first work that tackles the problem of adversarial attacks against FR systems in a CR scenario. For such a reason, in what follows, we briefly cite a few articles related to the topics of the cross-resolution FR and adversarial attacks against an FR system.

2.1. Cross-Resolution Face Recognition

CR scenarios are met whenever images at different resolutions have to be matched. Such a situation typically happens, for example, in biometric and forensics applications. Super-Resolution (SR) techniques are among the most studied solutions to such a problem, and Singh et al. (2018) proposed to synthesize high-resolution faces from low-resolution ones by employing a multi-level sparse representation of the given inputs. Zangeneh et al. (2020) formulated a mapping of the low- and the high-resolution images to a common space by leveraging a DL architecture made by two distinct branches, one for each image. Luo et al. (2019) exploited the dictionary learning approach based on learning multiple dictionaries, each being associated with a resolution. The most comprehensive study and widely tested method to improve an FR system's performance in a CR scenario was recently proposed by Massoli et al. (2020). In their work, the authors formulated a training procedure to fine-tune a state-of-the-art model to the CR domain. They tested their models on several benchmark datasets by showing their superior performance compared to the results available in the literature.

2.2. Face Recognition Adversarial Attacks

As we mentioned at the beginning of this section, we are the first to study adversarial attacks in a cross-resolution domain. Due to the lack of papers than can be directly compared to our study, in what follows we only briefly cite a few articles concerning adversarial attacks against FR systems. Sharif et al. (2016) demonstrated the feasibility and effectiveness of physical attacks by impersonating other identities using eyeglass frames with a malicious texture. Zhong and Deng (2020) observed the superior transferability properties of feature-based attacks compared to label-based ones. Moreover, they proposed a drop-out method for DCNNs to enhance further the transferability of the attacks. Song et al. (2018) proposed a three-player GAN architecture that leveraged a face recognition network as the third player in the competition between generator and discriminator. Dong et al. (2019) successfully performed black-box attacks on FR models and demonstrated their effectiveness in a real-world deployed system.

3. Adversarial Attacks

3.1. Carlini and Wagner

Carlini and Wagner (Carlini and Wagner, 2017) (CW) formulated one of the strongest currently available attacks. The CW-L_2 attack is formalized as:

$$\min \ c \cdot f(\tfrac{1}{2}\tanh(\mathbf{w}) + 1) + \| \tfrac{1}{2}(\tanh(\mathbf{w}) + 1) - \mathbf{x} \|_2^2,$$

where $f(\cdot)$ is the objective function, $\mathbf{x}$ is the input image, $\mathbf{w}$ is the adversarial example in the tanh space, and c is a positive constant which value is set by exploiting a binary search procedure.

3.2. Elastic Net Attack to DNNs

The Elastic Net Attack (Chen et al., 2018) (EAD), leverages the elastic-net regularization which is a well known technique in solving high-dimensional feature selection problems (Zou and Hastie, 2005). It is based on the objective proposed in Carlini and Wagner (2017) and it conceives the CW-L_2 attack as a special case. EAD is formulated as:

$$\min_{\mathbf{x}} \ c \cdot f(\mathbf{x}, t) + \beta \| \mathbf{x} - \mathbf{x}_0 \|_1 + \| \mathbf{x} - \mathbf{x}_0 \|_2^2,$$

where $f(\cdot)$ is the objective as in the CW-L_2 attack, t is the target class, $\mathbf{x}_0$ is the input image, t is the target label, $\mathbf{x}$ is the adversarial instance, c is a parameter found by binary search, and β represents the weight of the L_1 penalty term.

3.3. Jacobian Saliency Map Attack

The Jacobian Saliency Map Attack (Papernot et al., 2016) (JSMA) exploits an "input-perturbation-to-output" mapping. Differently from the backpropagation-based attacks, JSMA leverages the model derivative concerning the classification output rather than the derivative of the loss function. The attack is formalized as: $\arg\min_{\delta_{\mathbf{x}}} \parallel \delta_{\mathbf{x}} \parallel$ s.t. $\mathbf{F}(\mathbf{X} + \delta_{\mathbf{x}}) = \mathbf{Y}^{*}$, where $\mathbf{F}$ is the function learned by the DNN, $\mathbf{X}$ and $\mathbf{Y}^{*}$ are the input and output of the model, respectively, and $\delta_{\mathbf{x}}$ is the adversarial perturbation defined upon the evaluation of the model input saliency map.

3.4. Deep Representations Attacks

Differently from the previously mentioned attacks, the Deep Representations (Sabour et al., 2015) (DR) attack focuses on the manipulation of image features. It is formulated as an optimization problem which aims at finding the closest perturbed image, to the original one, whose descriptor is as close as possible to the one of a target image named the "guide image". Specifically, the adversarials crafting procedure is the following: $\mathbf{I}_{\alpha} = \arg\min_{\mathbf{I}} \parallel \phi_k(\mathbf{I}) - \phi_k(\mathbf{I}_g) \parallel_2^2$; subject to $\parallel \mathbf{I} - \mathbf{I}_s \parallel_{\infty} < \delta$, where $\phi(\cdot)_k$ is the descriptor extracted at layer k of the threatened model, $\mathbf{I}_s$ and $\mathbf{I}_g$ are the source and target images, respectively, $\mathbf{I}_{\alpha}$ is the adversarial example, and δ is he maximum allowed perturbation in terms of the L_{∞} norm.

4. Experimental Approach

4.1. Dataset and Models

In our experiments, we use the ~2.9M images shared among the 8631 identities contained in the training set of the VGG-Face2 (Cao et al., 2018) dataset. To construct the gallery and the queries, we divide the training set into two splits. Concerning the gallery, we evaluate a single template for each identity as the average features vector among all the corresponding face images. Regarding the queries, we randomly select 100 identities, and for each of them, we randomly pick ten correctly classified images, ending up with 1000 queries.

Concerning the learning models, we analyze the performance of two DCNNs: the face classifier from Cao et al. (2018) and the CR-trained one from Massoli et al. (2020). They share the same structure, i.e., a ResNet-50 (He et al., 2016) architecture equipped with Squeeze-and-Excitation (Hu et al., 2017) blocks. For both models, we adopt the same preprocessing steps for the images. First, following the same procedure as in Massoli et al. (2020), we synthesize different resolution versions of the input that allow us to evaluate the performance of the models in a cross-resolution scenario. Specifically, in our analysis, we consider images at 16, 24, 64, and 256 pixels (shortest side). Next, each image is resized to have the shortest side of 256 pixels, and then it is cropped to a square picture of size 224*x*224 pixels. Finally, we subtract the channel mean from each pixel.

4.2. Adversarial Attacks

Concerning the generation of the adversarial instances, we exploit the five algorithms we described in Section 3. We use the implementations available in the *foolbox* library (https://foolbox.readthedocs.io/en/stable/), with the only exception of the DR one that we build on top of the L-Broyden-Fletcher-Goldfarb-Shanno (L-BFGS) (Szegedy et al., 2013), optimization procedure. More precisely, the L-BFGS algorithm requires a function to optimize. To our aim, we implement such a function by employing a k-NN algorithm as guidance in the adversarial search. We fit the classifier to the gallery templates we mentioned at the beginning of this section. Then, we start the crafting procedure and stop it as soon as the k-NN classifies the malicious image as belonging to the targeted identity. In Figure 1, we report a schematic view of the procedure we just described.

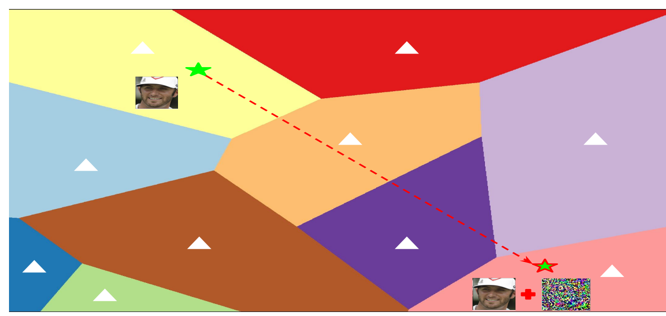

Figure 1: Schematic representation of our approach to crafting DR attacks. The colored regions are the k-NN decision boundaries for ten different identity templates (white triangles). The initial location of the green star represents a correctly classified features vector. The adversarial features vector's final position is represented by the red encircled star.

4.3. Face Identification Metrics

FR systems typically deal with sensitive scenarios such as biometric and forensics applications. Hence, different error types have distinct relevance while evaluating system performance, and a simple accuracy measure is not enough to properly evaluate and compare the performance of FR systems. Instead, as mentioned in Section 1, we focus our study on the face identification protocol. Specifically, we consider both the close- and open-set settings.

Concerning the close-set setting, we evaluate the Cumulative Match Characteristic (CMC), a metric that represents a summarized accuracy evaluated on mated searches only, i.e., considering queries that correspond to identities already available the gallery. The CMC value at rank one is usually named "hit rate," and it is the most typical summary indicator of an algorithm's efficacy. As we mentioned above, we select 100 identities to construct the queries. Thus, we end up with a gallery containing 8631 identities that comprise a hundred mated ones and 8531 un-mated ones acting as "distractors".

In the open-set setting, differently from the close-set one, we consider both mated and un-mated queries. To this aim, we remove half of the queries identities from the gallery, ending up with 50 mated and 50 un-mated persons and a gallery containing 8581 templates. With that set, there are two different types of errors that are usually evaluated, i.e., the False Positive Identification Rate (FPIR) and the False Negative Identification Rate (FNIR) or "miss rate". Concerning the former, it represents the number of un-mated queries that return a positive match at or above a specific similarity threshold. On the other hand, the FNIR represents the number of mated searches that return candidates with a similarity score below the threshold or outside the top R ranks.

The FNIR and FPIR, parametrized by the similarity threshold, can be combined to construct the Detection Error Tradeoff (DET), which is typically used to report the two types of error trade-off. We use the DET to evaluate the performance of the learning models in the experiments.

5. Experimental Results

We dedicate this section to report the results of our experimental campaigns. As we mentioned in Section 1, we aim to study the behavior of DL-based FR systems when threatened by adversarial attacks in a CR domain. Concerning the FR, as backbone features extractors, we consider the well-known DCNN from Cao et al. (2018) that set the state-of-the-art on the NIST datasets (Klare et al., 2015; Whitelam et al., 2017; Maze et al., 2018) and the CR model from Massoli et al. (2020) that set the state-of-the-art in the cross-resolution domain.

To craft adversarial examples, we harness the algorithms we described in Section 3. Moreover, being interested in the CR scenario, we consider input faces at 16, 24, 64, and 256 pixels (shortest side). Concerning the FR task, we keep the gallery at the original resolution.

As mentioned in Section 2, to our knowledge, we are the first to conduct this type of study. Thus, a direct comparison with previously published works is not possible. Hence, in what follows, we only report our results. We hope that our study will stimulate further researches in this direction. Throughout this section, we refer to the model from Cao et al. (2018) as "Base" model and to the one from Massoli et al. (2020) as "Cross-Resolution" model.

5.1. Threatening the Classification

We report the results from the attacks against the classification in Table 1. Concerning the attacks, we use the following configurations. For JSMA, we consider 1000 iterations, a perturbation per pixel equals to 0.1, 0.3, and 0.5 (percentage over the allowed pixel range), and a maximum number of times each pixel can be modified of 10. For CW-L_2, we consider 10 binary search steps and 10 and 100 iterations. Concerning EAD, we use the same parameters as for the CW-L_2 attack and a value for the weight of the L_1 penalty term equals to 0.1 and 1. Furthermore, since the DR (Sabour et al., 2015) attack is the least time demanding compared to the others, we enlarge the set of hyperparameters for it. Thus, we dedicate Figure 2 to report their results.

From Table 1, we notice that there is no clear signature for which model is more robust against adversarial attacks. On the other hand, we see that, on average, an adversary's success rate decreases as the resolution increases while keeping the attack configuration fixed. Let us now turn our attention to a single attack, for example, CW-L_2. It is interesting to notice the impact of a different choice of hyperparameters. Indeed, even though from the configuration (10-10), the "Base" model seems to be more resilient compared to the "Cross-Resolution" one, this is not true. Indeed, by just increasing the strength of the attack, i.e., (10-100) configuration for which we grow the number of steps, we reach 100% of attack success rate for both models.

From Figure 2 we observe that it is undeniable that the deep features extracted by the "Cross-Resolution" model are much more robust than those extracted from the "Base" NN. Thus, confirming our previous assertion about the benefit of CR training. From the first plot of Figure 2, we see that the success rate of the attack is almost 0% for the "Base" model. Instead, in the second plot, it looks like that both models have the same resilience. This is not in contrast with our previous conclusions. Indeed, as it has been shown in appendix 1 of Massoli et al. (2020), the "Base" model is not able to generate meaningful deep representation at very low resolutions. Thus, it is almost impossible to craft targeted attacks based on deep features. To sustain even more our assertion, we run a test with untargeted DR attacks in which we easily reach a success rate of 100% for

Table 1: Attack success rate against classification for "Base" and "Cross-Resolution" models. The first column reports the specific configuration used for each attack. The four values reported in the second and third main columns represent the success rate at a resolution of 16, 24, 64, and 256 pixels, respectively.

Attack Configuration	Attack Success Rate (%)							
	Base Model				Cross-Resolution Model			
	16	24	64	256	16	24	64	256
JSMA (1000-0.1-1.0)	76.1	61.8	25.5	11.5	65.5	62.8	17.1	6.9
JSMA (1000-0.3-1.0)	96.6	92.5	75.7	61.2	96.0	94.7	70.0	50.1
JSMA (1000-0.5-1.0)	98.5	95.8	86.4	76.6	97.6	97.0	**100.**	69.6
CW-L_2 (10-10)	82.9	72.9	45.9	32.7	86.4	83.3	52.8	37.4
CW-L_2 (10-100)	**100.**	**100.**	**100.**	**100.**	**100.**	**100.**	**100.**	**100.**
EAD (10-0.1-10)	95.7	98.2	94.5	87.0	96.7	99.6	98.8	98.5
EAD (10-0.1-100)	**100.**	**100.**	**100.**	**100.**	**100.**	**100.**	**100.**	**100.**
EAD (10-1.0-10)	83.4	85.1	50.2	27.9	72.6	94.4	86.9	73.8
EAD (10-1.0-100)	98.5	99.8	98.7	91.0	97.5	99.8	**100.**	99.6

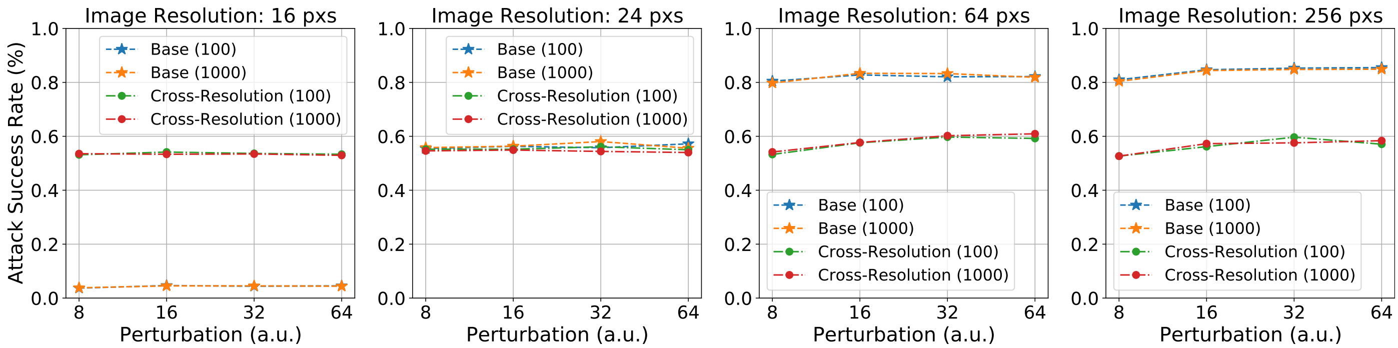

Figure 2: DR (Sabour et al., 2015) attack success rate as function of the maximum allowed perturbation δ considering 100 and 1000 iteration steps. Each plot represents a different input resolution.

the "Base" model.

Finally, we can notice that from our results, there is no clear evidence in favor of a specific metric since with the proper hyperparameters, we reached high success rates with the L_1, L_2, and L_∞.

5.2. *Threatening the Face Recognition*

We now turn our attention to DL-based FR systems. We begin our analysis by considering the face identification protocol in the close-set scenario, and we then move the open-set one. We refer the reader to Section 4 for a detailed description of the metrics we use to assess the performance of the systems under analysis.

5.2.1. *Close-set*

As mentioned in Section 4, we use the CMC to evaluate the performance of the threatened models in the close-set scenario. Specifically, we summarize our results in Table 2 by reporting the hit rate, i.e., the CMC value at a rank equals to one, with the exception of the DR (Sabour et al., 2015) attack to which we dedicate Figure 3. From a defensive point of view, the more resilient a model, the lower the hit rate, while from an attacker perspective, it is the other way round.

By looking at Table 2 and Figure 3 we can assert that the DR attack is much more effective in fooling a DL-based FR system than the classification-based ones with respect to any type of metric. From the attacker's point of view, this is a fundamental result. Indeed, by comparing the results from Table 1 and Table 2, we see that even though the attacks fool the classification, it is not guaranteed that they can evade a similarity-based system. Thus, deep representation attacks might be a better choice to attack an FR system. Moreover, we see how the "Cross-Resolution"-based system exhibits higher robustness than the one based on the "Base" model. Thus, again, we find that DCNNs benefit from a CR training approach (Massoli et al., 2020) in terms of resilience to adversarial attacks. Indeed, it is undeniable that the "Cross-Resolution"-based system is much more resilient against adversarial attacks than the "Base"-based one across all resolutions.

5.2.2. *Open-set*

To report the results for the face identification protocol in the open-set setting, we exploit the DET. Two fundamental aspects differentiate the DET from the CMC. Indeed, the former applies a threshold among the similarity of the features, and it comprises queries of identities that are not present in the gallery. Instead, the latter does not use any threshold, i.e., it does not discern among "weak" and "strong" similarity scores, and it requires queries related to already known identities.

As we mentioned in Section 4, the DET represents the er-

Table 2: Attacks hit rate. The first column reports the configuration for each attack. The four values reported in the second and third main columns are the results at a resolution of 16, 24, 64, and 256 pixels, respectively As a reference, we report in the first row the hit rate for the authentic images.

Attack Configuration	Hit Rate (%)							
	Base Model				Cross-Resolution Model			
	16	24	64	256	16	24	64	256
Auth	79.5	95.3	99.8	99.9	96.7	98.8	99.4	99.7
JSMA (1000-0.1-1.0)	12.1	10.7	12.9	12.2	11.9	9.8	9.4	13.0
JSMA (1000-0.3-1.0)	14.0	9.3	10.7	10.6	9.8	10.0	7.4	8.9
JSMA (1000-0.5-1.0)	13.6	10.6	10.0	10.3	10.0	10.2	3.0	6.8
CW-L_2 (10-10)	10.9	6.5	6.1	3.7	10.8	9.3	5.5	5.1
CW-L_2 (10-100)	7.6	4.1	6.1	2.3	9.2	9.3	3.6	4.6
EAD (10-0.1-10)	31.8	32.6	**27.8**	25.1	19.2	16.8	19.4	19.7
EAD (10-0.1-100)	17.5	9.7	6.3	6.2	13.8	11.6	6.8	5.3
EAD (10-1.0-10)	**44.8**	**38.0**	26.7	**25.5**	**20.8**	**25.7**	**20.1**	**21.7**
EAD (10-1.0-100)	34.8	30.3	20.7	16.8	17.3	16.5	17.4	17.2

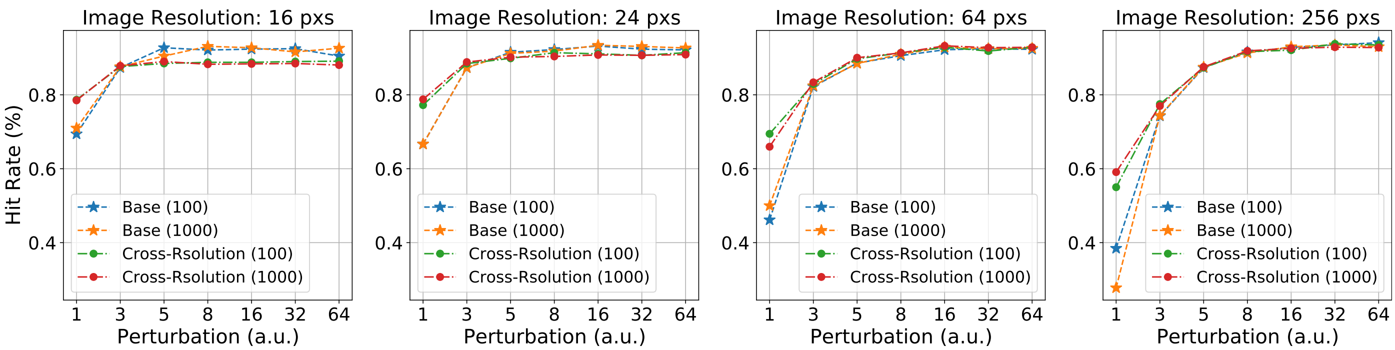

Figure 3: DR (Sabour et al., 2015) hit rate as function of the maximum allowed perturbation δ considering 100 and 1000 attack steps. Each plot represents a different input resolution.

ror trade-off between the FNIR and the FPIR. To summarize the performance of the FR systems, we report the FPIR at a reference value of the FNIR equals to $1.e^{-2}$. Compared to the close-set settings, the adversary's goal is to lower the curve as much as possible, while from a defensive point of view, a higher curve represents a more resilient model. The results are reported in Table 3 with the exception of DR (Sabour et al., 2015) to which we dedicate Figure 4.

Analyzing the results reported in Table 3 and Figure 4 we obtain the same conclusions we report for the close-set setting. Specifically, by comparing the results from Table 3 to the ones in Figure 4 we see that the DR attack is much more effective in fooling the FR system compared to others and that the "Cross-Resolution"-based system is much more resilient than the "Base"-based one against adversarial attacks.

6. Conclusions

DCNN-based FR systems leverage the representation power of learning models. Unfortunately, they also share their weaknesses. Indeed, it has been recently shown that these systems suffer a drastic drop in their performance when tested in a cross-resolution domain. The situation becomes even worse when an adversary comes into play. Indeed, an FR system can be deceived by adversarial examples. These weaknesses pose a severe limit to the spread of these systems to sensitive real-world applications such as biometric systems and forensics.

In such a context, we proposed our analysis in which we compared the resilience to adversarial attacks of FR systems based on the deep features extracted by NNs in a CR scenario. We studied two different DCNN models: a former one, trained only on high-resolution images and a latter one, trained on a cross-resolution domain. To generate adversarial instances, we harnessed several algorithms based on different metrics and objectives, and we craft malicious samples considering input images at a resolution of 16, 24, 64, and 256 pixels. Concerning the measures of the performance of the FR systems, we adopted the face identification protocol. Specifically, we considered the close- and open-set settings for which we evaluated the CMC and DET.

From our analysis, we notice that, given a specific configuration, the attack success rate is higher at lower resolutions, for example, at 16 and 24 pixels, than at higher ones, such as 64 and 256 pixels. Such behavior was somehow expected since, at a very low-resolution part of the face information can be lost, thus simplifying the effort of an adversary.

By looking at the results from the FR systems, it is evident that a DCNN benefits from a CR training procedure since it empowers the learning model to extract more robust deep representations. Moreover, we observed that DR attacks represent

Table 3: FPIR@FNIR=$1.e^{-2}$. The first column reports the configuration for each attack. The four values reported in the second and third main columns are the results at a resolution of 16, 24, 64, and 256 pixels, respectively. As a reference, we report in the first row the results for the authentic images.

Attack Configuration	FPIR@FNIR=$1.e^{-2}$							
	Base Model				Cross-Resolution Model			
	16	24	64	256	16	24	64	256
Auth	75.0	40.8	0.8	1.0	38.6	20.2	3.6	3.2
JSMA (1000-0.1-1.0)	99.3	99.1	100.	95.1	99.1	98.4	100.	98.1
JSMA (1000-0.3-1.0)	99.0	99.1	97.2	99.7	97.8	98.6	99.0	100.
JSMA (1000-0.5-1.0)	98.0	98.1	98.2	97.0	99.4	98.6	99.0	98.7
CW-L_2 (10-10)	99.5	98.1	99.5	97.4	99.0	**98.1**	98.9	98.9
CW-L_2 (10-100)	100.	99.0	99.5	99.4	99.6	**98.1**	99.6	99.2
EAD (10-0.1-10)	**95.3**	**93.2**	98.7	99.5	98.4	98.8	**96.0**	97.6
EAD (10-0.1-100)	98.0	99.4	99.4	99.0	100.	98.8	98.6	99.2
EAD (10-1.0-10)	95.6	96.3	98.3	**95.3**	**96.3**	**98.1**	96.7	**96.7**
EAD (10-1.0-100)	98.8	97.9	**97.1**	98.6	98.6	**98.1**	99.0	97.7

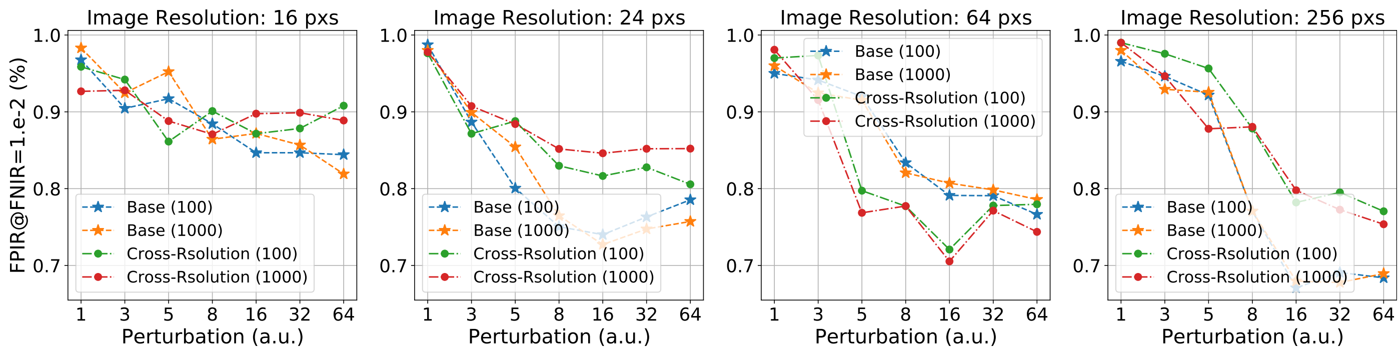

Figure 4: FPIR@FNIR=1.e^{-2} for the DR (Sabour et al., 2015) attack as function of the maximum allowed perturbation δ considering 100 and 1000 attack steps. Each plot represents a different input resolution.

a much greater menace to an FR system than the ones based on the classification output of the threatened models for each of the considered metrics, i.e., L_1, L_2 and L_∞. Such a result held for the close- as well as for the open-set settings.

References

Ahonen, T., Hadid, A., Pietikainen, M., 2006. Face description with local binary patterns: Application to face recognition. IEEE TPAMI , 2037–2041.

Amato, G., Falchi, F., Gennaro, C., Massoli, F.V., Passalis, N., Tefas, A., Trivilini, A., Vairo, C., 2019. Face verification and recognition for digital forensics and information security, in: ISDFS, IEEE. pp. 1–6.

Biggio, B., Corona, I., Maiorca, D., Nelson, B., Šrndić, N., Laskov, P., Giacinto, G., Roli, F., 2013. Evasion attacks against machine learning at test time, in: ECML PKDD, Springer. pp. 387–402.

Cao, Q., Shen, L., Xie, W., Parkhi, O.M., Zisserman, A., 2018. Vggface2: A dataset for recognising faces across pose and age, in: International Conference on Automatic Face & Gesture Recognition, IEEE. pp. 67–74.

Carlini, N., Wagner, D., 2017. Towards evaluating the robustness of neural networks, in: Symposium on security and privacy, IEEE. pp. 39–57.

Chen, P.Y., Sharma, Y., Zhang, H., Yi, J., Hsieh, C.J., 2018. Ead: elastic-net attacks to deep neural networks via adversarial examples, in: Thirty-second AAAI conference on artificial intelligence.

Cheng, Z., Zhu, X., Gong, S., 2018. Surveillance face recognition challenge. arXiv preprint arXiv:1804.09691 .

Deng, J., Guo, J., Xue, N., Zafeiriou, S., 2019. Arcface: Additive angular margin loss for deep face recognition, in: CVPR, IEEE. pp. 4690–4699.

Dong, Y., Su, H., Wu, B., Li, Z., Liu, W., Zhang, T., Zhu, J., 2019. Efficient decision-based black-box adversarial attacks on face recognition, in: CVPR, IEEE. pp. 7714–7722.

Ekenel, H.K., Sankur, B., 2005. Multiresolution face recognition. Image and Vision Computing 23, 469–477.

Goswami, G., Ratha, N., Agarwal, A., Singh, R., Vatsa, M., 2018. Unravelling robustness of deep learning based face recognition against adversarial attacks, in: Thirty-Second AAAI Conference on Artificial Intelligence.

Grother, P., Grother, P., Ngan, M., Hanaoka, K., 2019. Face Recognition Vendor Test (FRVT) Part 2: Identification. US Department of Commerce, National Institute of Standards and Technology.

He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep residual learning for image recognition, in: CVPR, IEEE. pp. 770–778.

Hu, J., Shen, L., Sun, G., 2017. Squeeze-and-excitation networks. arxiv.

Kakizaki, K., Yoshida, K., 2019. Adversarial image translation: Unrestricted adversarial examples in face recognition systems. arXiv:1905.03421.

Klare, B.F., Klein, B., Taborsky, E., Blanton, A., Cheney, J., Allen, K., Grother, P., Mah, A., Jain, A.K., 2015. Pushing the frontiers of unconstrained face detection and recognition: Iarpa janus benchmark a, in: CVPR, IEEE. pp. 1931–1939.

Kurakin, A., Goodfellow, I., Bengio, S., 2016. Adversarial examples in the physical world. arXiv preprint arXiv:1607.02533 .

LeCun, Y., Bengio, Y., Hinton, G., 2015. Deep learning. nature 521, 436–444.

Liu, C., Wechsler, H., 2002. Gabor feature based classification using the enhanced fisher linear discriminant model for face recognition. IEEE Transactions on Image processing 11, 467–476.

Luo, X., Xu, Y., Yang, J., 2019. Multi-resolution dictionary learning for face recognition. Pattern Recognition 93, 283–292.

Massoli, F.V., Amato, G., Falchi, F., 2020. Cross-resolution learning for face recognition. Image and Vision Computing , 103927.

Massoli, F.V., Amato, G., Falchi, F., Gennaro, C., Vairo, C., 2019. Improving multi-scale face recognition using vggface2, in: International Conference on Image Analysis and Processing, Springer. pp. 21–29.

Maze, B., Adams, J., Duncan, J.A., Kalka, N., Miller, T., Otto, C., Jain, A.K., Niggel, W.T., Anderson, J., Cheney, J., et al., 2018. Iarpa janus benchmark-c: Face dataset and protocol, in: ICB, IEEE. pp. 158–165.

Papernot, N., McDaniel, P., Jha, S., Fredrikson, M., Celik, Z.B., Swami, A., 2016. The limitations of deep learning in adversarial settings, in: 2016 IEEE European symposium on security and privacy, IEEE. pp. 372–387.

Qiu, H., Xiao, C., Yang, L., Yan, X., Lee, H., Li, B., 2019. Semanticadv: Generating adversarial examples via attribute-conditional image editing. arXiv preprint arXiv:1906.07927 .

Sabour, S., Cao, Y., Faghri, F., Fleet, D.J., 2015. Adversarial manipulation of deep representations. arXiv preprint arXiv:1511.05122 .

Sharif, M., Bhagavatula, S., Bauer, L., Reiter, M.K., 2016. Accessorize to a crime: Real and stealthy attacks on state-of-the-art face recognition, in: SIGSAC CCS, ACM. pp. 1528–1540.

Singh, M., Nagpal, S., Singh, R., Vatsa, M., Majumdar, A., 2018. Magnifyme: Aiding cross resolution face recognition via identity aware synthesis. arXiv preprint arXiv:1802.08057 .

Song, Q., Wu, Y., Yang, L., 2018. Attacks on state-of-the-art face recognition using attentional adversarial attack generative network. arXiv preprint arXiv:1811.12026 .

Spaun, N.A., 2011. Face recognition in forensic science, in: Handbook of face recognition. Springer, pp. 655–670.

Sundararajan, K., Woodard, D.L., 2018. Deep learning for biometrics: a survey. ACM Computing Surveys (CSUR) 51, 65.

Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., Fergus, R., 2013. Intriguing properties of neural networks. arXiv preprint arXiv:1312.6199 .

Turk, M.A., Pentland, A.P., 1991. Face recognition using eigenfaces, in: Proceedings. 1991 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, IEEE. pp. 586–591.

Wang, M., Deng, W., 2018. Deep face recognition: A survey. arXiv preprint arXiv:1804.06655 .

Whitelam, C., Taborsky, E., Blanton, A., Maze, B., Adams, J., Miller, T., Kalka, N., Jain, A.K., Duncan, J.A., Allen, K., et al., 2017. Iarpa janus benchmark-b face dataset, in: CVPR Workshops, IEEE. pp. 90–98.

Zangeneh, E., Rahmati, M., Mohsenzadeh, Y., 2020. Low resolution face recognition using a two-branch deep convolutional neural network architecture. Expert Systems with Applications 139, 112854.

Zhang, K., Zhang, Z., Cheng, C.W., Hsu, W.H., Qiao, Y., Liu, W., Zhang, T., 2018. Super-identity convolutional neural network for face hallucination, in: European Conference on Computer Vision (ECCV), pp. 183–198.

Zhong, Y., Deng, W., 2020. Towards transferable adversarial attack against deep face recognition. arXiv preprint arXiv:2004.05790 .

Zou, H., Hastie, T., 2005. Regularization and variable selection via the elastic net. Journal of the royal statistical society: series B 67, 301–320.

Zou, W.W., Yuen, P.C., 2011. Very low resolution face recognition problem. IEEE Transactions on image processing 21, 327–340.