Saliency-Guided Textured Contact Lens-Aware Iris Recognition

Lucas Parzianello
University of Notre Dame
lparzianello@nd.edu

Adam Czajka
University of Notre Dame
aczajka@nd.edu

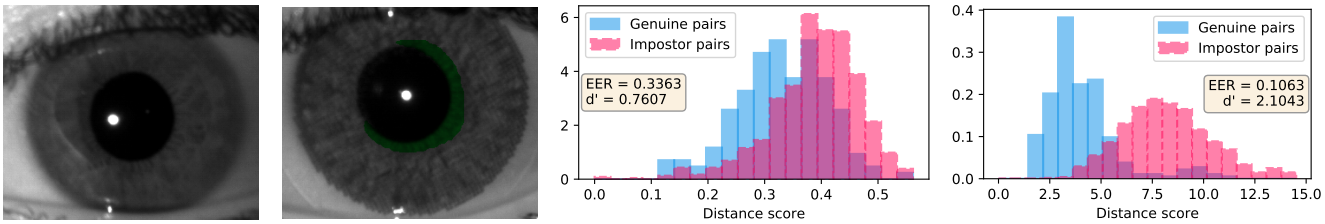

(a) (b) (c) (d)

Figure 1: An example live iris without (a) and with (b) a textured contact lens, leaving only a small usable area highlighted in green on the right picture. Comparisons between images such as (a) and (b), done with a traditional iris recognition method, end up with a high probability of false non-match for reasonable acceptance threshold (c). The method proposed in this paper offers much better separation between genuine and impostor scores, and thus much higher probability of correct matching between an iris covered by a textured contact lens and a live iris, as shown in histogram (d).

Abstract

Iris recognition requires an adequate level of the iris texture being visible to perform a reliable matching. In case when a textured contact lens covers the iris, a false non-match is reported or a presentation attack is detected. There are, however, scenarios in which one wants to maximize the probability of a correct match despite the iris texture being being partially or mostly obscured, for instance when a non-cooperative subject conceals their identity by purposely wearing textured contact lenses. This paper proposes an iris recognition method designed to detect and match portions of live iris tissue still visible when a person wears textured contact lenses. The proposed method includes (a) a convolutional neural network-based segmenter detecting partial live iris patterns, and (b) a Siamese network-based feature extraction model, trained in a novel way with images having non-iris information removed by blurring, to guide the network towards salient live iris features. Experiments matching pairs of iris images in which the iris is not wearing a lens in one image and is wearing a textured contact lens in the other, show a lower EER=10.6% for the proposed algorithm, compared to state-of-the-art iris code-based iris recognition (EER=33.6%). The source codes of the method are offered along with the paper.

1. Introduction

Iris recognition is regarded as one of the most secure and widely used biometric authentication methods. It has been employed in large-scale applications such as national identification [7], as well as an authentication factor for high-security facilities. However, current models for iris recognition may take subject cooperation as granted, which might not be always the case, particularly for large-scale and lightly supervised applications. Works in Presentation Attack Detection (PAD) [8, 6, 11] show that it is possible, especially if the attack types are known, to identify with high accuracy whether a given iris image has a genuine origin, or if it was altered in some form: either digitally or by presenting a physical imitation of the iris. The latter includes subjects using textured contact lenses, either for deceitful or cosmetic purposes, which in a state-of-the-art iris recognition system would result in a false non-match when compared to the unobstructed iris of the same person.

The question arises: instead of rejecting an image of the iris covered by a textured contact lens, can we build a method that generates a correct match despite severe occlusions? Or at least a system that increases the probability of a correct match in such case, compared to the state-of-the-art?

We made an observation that the opening in the textured contact lenses, allowing the light to pass through it and reach the retina, is larger that the actual pupil in most of the contact lenses, as highlighted in Fig. 1. This is to account for different pupil sizes and dilation, and leaves a tiny region of authentic iris texture still exposed. This paper proposes a method which takes advantage of this "leftover" genuine information in order to perform the iris recognition. We show that in the majority of cases, despite the subjects wearing textured lenses that would deceive a traditional method, the proposed solution shows a significant separation between genuine and impostor distributions: the d' of 0.76 using a state-of-the-art iris matcher, versus the d' of 2.10 using the proposed TCLA matcher on an unseen dataset. This distance increase represents a practical reduction in the Equal Error Rate (EER), from 33% to 10%. While these recognition results are certainly worse than those observed when ample iris area is available for matching, the probability of getting a correct match for a hypothesized person concealing their identity is significantly higher using the proposed method, when compared to a state-of-the-art iris recognition approach.

There are three contributions offered by this paper. The **first contribution** is a method for **Textured Contact Lens-Aware (TCLA) iris segmentation**, which is applied after the sample is classified as belonging to a subject wearing a textured contact lens. This method locates the iris image regions not affected by the textured contact lens pattern, as opposed to traditional methods locating the entire iris annulus, regardless of authenticity of the pattern inside the annulus. The segmenter is based on Mask R-CNN architecture and was fine-tuned with manually segmented iris samples with textured contact lenses placed on the cornea.

Moving towards our goal of recognition, our **second contribution** is the **TCLA matcher**: an algorithm based on a Siamese Neural Network architecture to match iris images of unknown subjects wearing textured contact lenses against irises enrolled without contact lenses. This TCLA network is trained in a novel way, in which non-salient image regions, corresponding to non-iris areas, were suppressed by low-pass filtering applied to training data, following [5]. This model thus extracts the best encoding, given the tiny iris information available, by far outperforming the recognition performance achieved by a state-of-the-art human-driven BSIF-based method employing the same iris segmenter to mask out the contact lenses, even after selection of optimal BSIF encoding filters to take these occlusions into account. This shows that a simple adaptation of state-of-the-art method to get better encoding may not be sufficient.

As our **third contribution**, we make the manually segmented iris masks prepared for occluded irises available along with the paper. The annotations mark the areas that are uncovered by the contact lenses and can facilitate future attempts at maximizing recognition performance in situations when subjects make an attempt of hiding their identity by wearing off-the-shelf textured contact lenses. The annotations correspond to images taken from the publicly available *NDIris3D* dataset [16], and is composed of 1171 iris images captured from 119 subjects, all wearing regular or irregular textured contact lenses. (The *NDIris3D* dataset has also corresponding images free from contact lens occlusions, captured from the same subjects.)

This is the first attempt, known to us, to offer a method that increases the probability of correct iris-based matching in case of subjects wearing textured contact lenses, and having their unaltered images previously enrolled. This work should contribute to a better detection of those using textured contact lenses with an intention to hide their identity and not be recognized by an iris recognition system, *e.g.* at a border checkpoint.

2. Related Work

This section provides a concise overview of works related to iris recognition in situations when limited information about the iris texture is available.

Yadav *et al.* [21] explored the effects of textured contact lenses for iris recognition and contribute with a database of irises wearing a variety of contact lens types (soft and textured), lens makers and models, and iris sensors. Similarly, Fang *et al.* [15] reported a significant drop in performance when live irises are compared to their corresponding samples with textured contact lenses present on the cornea. None of these works proposed methods to increase the matching performance when one of the irises being compared is covered by textured contact lens.

Rathgeb *et al.* [19] propose to reduce the iris recognition algorithm complexity by using what the authors called *reliability masks*. In their experiments using only 5% of the iris code bits ended up with an EER below 5%. This supports the claim that it is possible to extract sufficient information for matching by using a small fraction of the iris. However, since the authors are not making the comparison with irises wearing textured contact lenses, this selected fraction of the iris code is composed by the most reliable bits of the code. This is something that cannot be assumed in the problem we consider in this paper: the information available after covering the iris with textured contact lenses may not be as significant for matching as the best-selected regions.

Hsieh *et al.* [18] were successful in developing a solution that distinguishes natural iris patterns from textured contact lens patterns, which resulted in a reduction in false match rate from over 10% to below 0.6% on an evaluation set of 200 enrolled and 200 probe iris images. After mask-

ing out the areas with a textured contact lens, their iris code generation is based on the Log-Gabor filter for feature extraction and the Hamming distance is used to compare the codes. To disentangle the authentic and contact lens textures, the authors designed a specialized hardware with a dual-band camera, and assumed that the iris texture and textured contact lens patterns are statistically independent in the spectral domain. This hardware and the assumptions allowed to use Independent Component Analysis in separating both textures. Our paper differs from that work in a sense that our solution is applicable to near infrared iris images compliant to ISO/IEC 19794-6, and this does not rely on non-standard or additional hardware compared to the existing commercial iris sensors, and thus can be added to the existing recognition pipelines with no additional hardware modifications. While their solution is agnostic to the iris pattern, it does not account for multi-colored textured contact lenses because of limitations in the spectral analysis, and as a result, their experiments only focused on single color textured contact lenses.

Concluding, we are not aware of previous works proposing iris recognition methods specialized to ISO-compliant NIR iris images, in a scenario when at least one iris in the comparison pair is covered by a textured contact lens.

3. Databases

3.1. Database for Fractional Iris Segmentation

For training the new iris segmenter, detecting small portions of iris texture in presence of textured contact lenses, the publicly-available *Notre Dame Photometric Stereo Iris Dataset (NDPSID)* [9] containing 5796 NIR iris images acquired by the LG IrisAccess 4000 sensor from 119 subjects was used. The subset utilized in the training and evaluation of the TCLA segmenter contains 1171 irises wearing textured contact lenses with both regular and irregular patterns. This portion of the data was manually annotated to indicate the regions that are unaffected by the textured contact lens, as shown in Figure 1(b). Manual annotations prepared for this work are made available along with this paper. This dataset was split into subject-disjoint training (60% of data) and validation (40% of data) subsets.

3.2. Database for Feature Extraction and Matching

The dataset described in 3.1 and used for the segmentation task was not suitable for the recognition tasks, as its number of samples is heavily biased to a total of less than 10 subjects, despite having over 100 subjects in the dataset. Also, the number of subjects wearing textured contact lenses is less than 10 in this benchmark. For this reason, we used the *NDIris3D* dataset, which contains iris images captured from 89 subjects with and without textured contact lenses, collected from two sensors: AD100 (5,237) images, and the LG4000 (3,488 images), making up a total of 8,725 iris images. After de-duplication, filtering, subject-disjoint split and pairing of images between genuine and impostor, the number of images pairs obtained is detailed in Table 1.

Table 1: Subject-disjoint partitioning of the *NDIris3D* dataset used in designing the TCLA matcher.

	Subjects		**Pairs**		
	#	**%**	**Genuine**	**Impostor**	**Total**
Training	53	60%	761	13,738	**14,499**
Validation	18	20%	263	1,620	**1,883**
Testing	18	20%	263	1,390	**1,653**
Total	**89**	**100%**	**1,287**	**16,748**	**18,035**

3.3. Data Cleaning and Partitioning

With the final goal of iris matching in mind, the data was split between enrolled-genuine and enrolled-impostor pairs. Detailed below is each image group and how they were extracted from the *NDIris3D* dataset.

Enrolled (gallery) images are assumed to be captured under ideal circumstances: they do not wear any contact lenses and, therefore, do not require TCLA segmentation. They may have different illumination properties and be captured using either one of the two sensors (AD100 or LG4000) represented in the data.

Genuine probe images come from the same eye as their enrolled counterpart. They wear any kind of textured contact lens pattern (regular or irregular) of any brand or color present in the dataset. They may also have been captured under any combination of illumination / sensor, which can differ from their enrolled counterpart.

Impostor probe images follow the same properties as genuine images, except that they come from different subjects than their enrolled pairs.

It is worth noticing that the same image wearing a textured contact lens can be either genuine or impostor, depending on to which enrolled subject it is being matched against. Hence, we will refer to an unknown iris that can be either genuine or impostor as a *probe* image / sample / iris. Moreover, instead of mentioning the number of individual images, it is more appropriate to quantify this data observing the number of pairs formed. In total, after filtering individual images we formed a sum of 1,287 genuine and 16,748 impostor pairs across 89 distinct subjects, as detailed by Table 1. The train, validation and test partitions are subject-disjoint, that is, all samples of a particular subject in the testing set remain unseen during the method's design phase.

4. The Proposed Method

4.1. Overview

Our method for fractional iris image matching has two components specialized to this scenario: the Textured Contact Lens-Aware (TCLA) segmenter and the TCLA matcher, as detailed in Figure 2. The deep learning-based segmenter detects iris regions not altered by the textured contact lenses. Both cropped iris images and normalized iris images (as suggested by early versions of the ISO/IEC 19794-6 and early iris recognition papers [13]) are next considered as plausible inputs to a Siamese network based matcher, returning the similarity score for detected fractions of iris images. Both components are trained in a way to be textured contact lens aware (hence the TCLA acronym), with details of this training provided in the following sections.

4.2. The TCLA Segmenter

Detectron2 [20], an open-source library for object detection and segmentation, has been used in this project to facilitate fine-tuning of the Mask R-CNN [17] model. We use Mask R-CNN as a backbone in order to make the segmentation task with limited number of samples feasible. First, we fine-tune it to complete the general task of iris segmentation without regards to the presence or the type of contact lens, later referred to as *regular segmentation*. Then we save the trained parameters and further fine-tune them to create a second model to detect authentically-looking parts of the iris, later referred to as *authentic segmentation*. With both segmenters, regular and authentic, trained using the same architecture and backbone, we are able to more reliably compare the performance of the iris matcher.

Traditional non-deep-learning iris segmenters, inspired on Daugman's seminal works, such as [3] may not be suitable for the task of distinguishing iris texture and contact lens texture and will, in most cases, classify areas belonging to the contact lens as live iris texture, unless the method is modified to fit the requirements of this particular task. Starting by applying the Mask R-CNN based segmenter on irises wearing textured contact lenses and a regular SegNet-based [1] iris segmenter on the enrolled (not occluded by lenses) images, we measured the the percentage of usable pixels in the iris region after the segmentation. The usable regions are composed by pixels qualified to be used in the matching. Figure 3 shows the distributions of these areas after the segmentation and normalization are applied on the images from the testing set. Both types of probes have a similar distribution, with median of usable iris area at around 12% and almost a fifth of the testing set showing less than 5% of usable area. Note that ISO/IEC 29794-6 recommends at least 70% of the iris annulus to be not occluded to classify the iris sample as "high utility" sample. Having around 10% and less information demonstrates how challenging this task may be for classical iris matchers, especially making strong assumptions on compliance of input data with ISO standards.

4.3. Embedding Saliency Into Training Data

The TCLA matcher utilizes the Siamese Neural Networks (SNN) architecture, with the ResNet-18 architecture as backbone, trained in a novel and non-standard way. In addition to regular data augmentations, the core component of training is "removing" non-iris information from training images and thus guiding the network towards salient regions (*i.e.* not occluded iris portions). This is done by either blurring regions excluded by the TCLA segmenter, or replacing them with random noise, as shown in Fig. 5. This process incorporates saliency identified by the segmenter to emphasize the properties that should be used by the network in the matching, while dampening the effects of the textured contact lenses. The reason of not applying the binary mask directly on the image is to avoid the introduction of border artifacts that have a potential of misleading the network into learning to match contours instead of iris pattern. This tendency of convolutional neural networks of relying on borders as opposed to texture is one of the challenges we intend to overcome by exploring different ways of guiding the network. A similar guidance approach in biometrics showed promising results in increasing the generalization of solutions for iris PAD [4].

Focusing now on Figure 5, we have the original iris pattern on the left and two additional columns of different types of saliency maps: Gaussian blur and random noise. These three are combined with two cropping and normalization methods making six types of images that were later evaluated. The Gaussian blur and random noise used the binary masks as guidance to be applied on the original images in a way to avoid passing along the sharp edges of the binary mask: the resulting image is the outcome of the application of a low-pass filter on the binary mask and using that as weights for the blend. The result is a final image that preserves the texture of interest, but also smoothly "removes" the non-iris areas.

The top row of Figure 5 shows the rubbersheet-based normalization of irises, followed by a folding process so that the pupillary boundary is in the center of the image, and the limbic boundary the top and bottom. This was done as a result of an implementation constraint: with the default normalized image being less than 224 pixels of height (64 pixels), either the network architecture or the input shape would need to be changed, and we opted for the latter to make this solution applicable to pre-trained backbones, usually built to accept 224×224 images. The concept presented in this paper, however, is in no way associated or dependent on the model's input resolution (within a reasonable range guarantying preservation of iris fea-

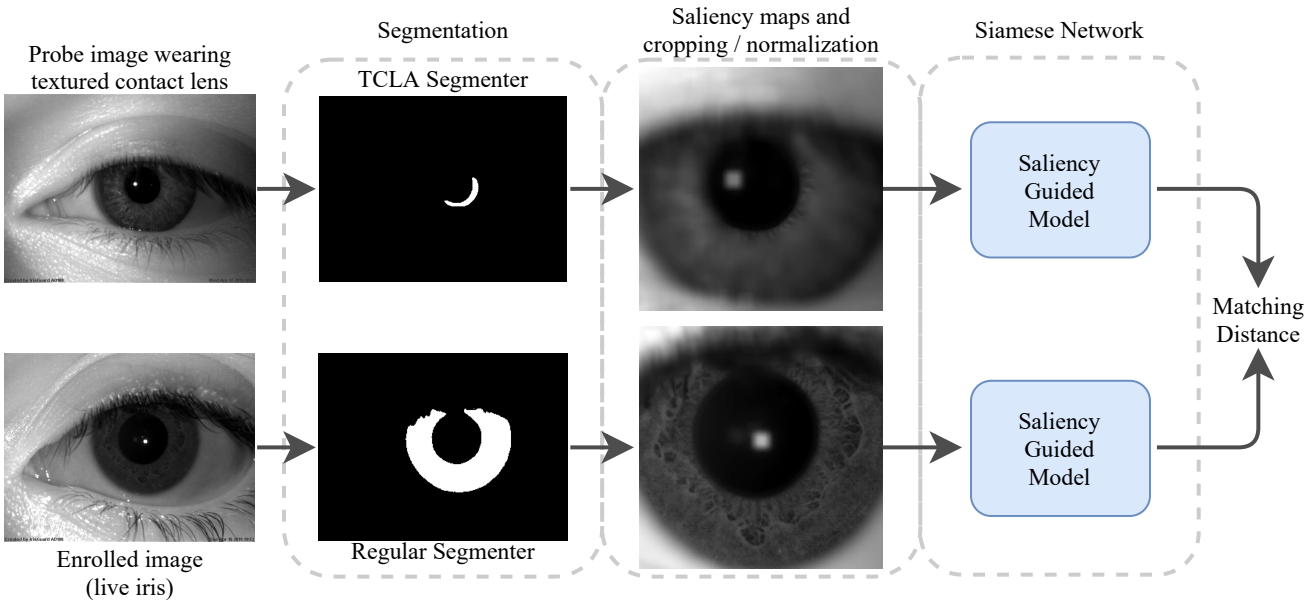

Figure 2: Overview of TCLA matcher; example using iris cropping and Gaussian blur as a saliency map.

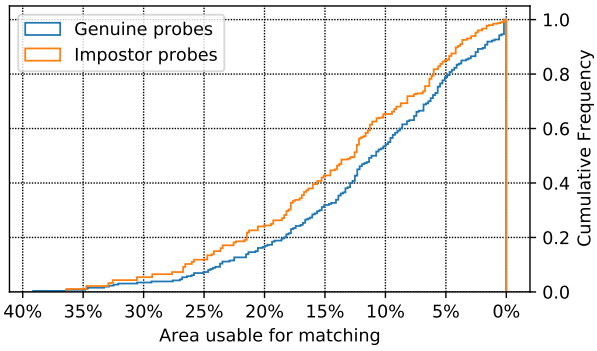

Figure 3: Percentage of authentic iris texture predicted to be authentic in each probe image from the testing set.

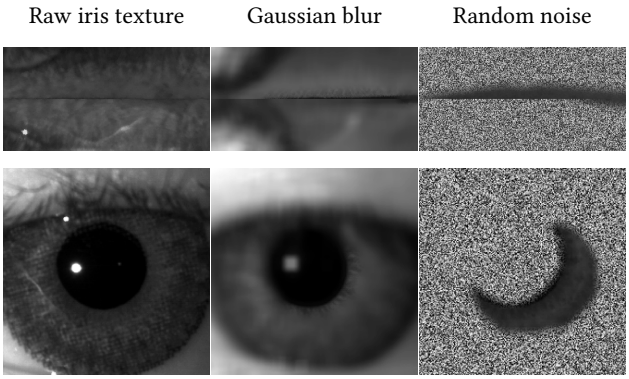

Figure 5: Examples of different approaches to saliency-based image data modification. All samples here are wearing a textured contact lens and do not necessarily belong to the same subject. The **second column** shows samples with non-iris regions were blurred. The **third column** shows samples with non-iris regions replaced with random noise. The **first row** shows normalized irises as input, while the **second row** shows center-cropped samples.

tures).

4.4. The TCLA Saliency-Guided Matching

In a SNN the weights are shared to process a pair of inputs independently, thus creating a pair of feature vectors, which are then compared using a distance function. These distances can be used in the loss according to the ground-truth of the original image pair in question (genuine or impostor in this case). As the goal of this matching network is to return a continuous comparison score given a pair or images as input, this architecture was a fitting choice for our goals. The TCLA matcher uses the same weights for both inputs, running a forward pass individually in the pre-processed images and computing the L2 distance of the resulting network embeddings. As in a traditional classifier, the final distances can be thresholded for a binary decision after the training is complete, or be used as a proxy for a confidence metric.

Preliminary experiments with a contrastive loss did not give good results, directing us to use a loss that could decrease the distance of genuine pairs and at the same time increase the distance between an impostor pair for the

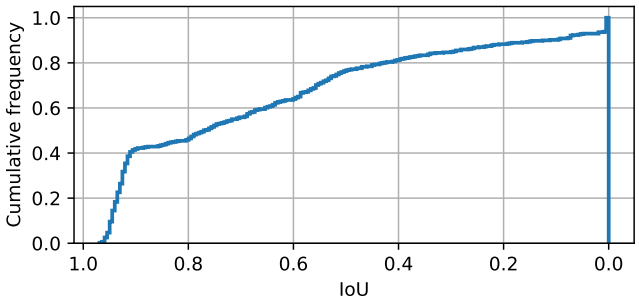

Figure 6: Cumulative Intersection over Union (IoU) obtained for the TCLA segmenter on the validation set. All irises wear a textured contact lens.

same enrolled image. For this reason, the following triplet loss $\mathcal{L}$ seemed a fit candidate to this problem:

$$\mathcal{L}\left(x_a, x_p, x_n\right) = \max\left(\left\|f\left(x_a\right) - f\left(x_p\right)\right\|^2 - \left\|f\left(x_a\right) - f\left(x_n\right)\right\|^2 + \alpha, 0\right)$$

where $f(x)$ is a network embedding for a sample x, x_a is the anchor (enrolled iris), x_p is a positive sample (genuine iris, belonging to the same class subject/eye), x_n is the negative sample (belonging to any other subject present in the same data partition: training or validation), and $\alpha = 1$ represents a margin and penalizes triplets for which the distance between the anchor and the positive sample is greater than the anchor and the negative sample plus the margin.

5. Experiments and Results

5.1. TCLA Segmenter

The sensitivity index (d') was used for comparisons among various scenarios. The d' is a statistical measure of separation between two distributions:

$$d' = \frac{|\mu_I - \mu_G|}{\sqrt{\frac{1}{2}\left(\sigma_I^2 + \sigma_G^2\right)}}$$

where μ_I, σ_I, μ_G, and σ_G are mean values and standard deviations of the impostor (I) and genuine (G) score distributions, respectively, and $|x|$ is the absolute value of x.

In order to assess the performance of the TCLA segmenter, we performed the evaluation on the left-out validation portion of the *NDPSID* dataset with ground-truth, corresponding to 496 labeled images. On this left-out data, the segmenter achieved a mean Intersection over Union (IoU) between the prediction and ground-truth of 0.669, and the cumulative distribution of IoU is illustrated in Figure 6. As it can be seen, over 40% of the images had an IoU of 0.9 or above, with around 75% of the images showing an IoU above or equal to 0.5. The figure also shows a non-negligible portion of images with IoU close to zero. These might be samples for which the network completely missed the "authentic" patches, or samples for which the textured contact lens covered the entire iris, thus there were no regions usable for a match.

Table 2: Mean values and units standard deviation of d' in five statistically independent inference experiments on the validation set with ("Random noise" and "Gaussian blur") and without saliency-guided training.

	No saliency	**Gaussian blur**	**Random noise**
Normalized	1.325 (±0.146)	**2.494** (±**0.234**)	0.459 (±0.160)
Cropped	1.465 (±0.150)	1.797 (±0.476)	0.816(±0.182)

5.2. TCLA Matcher

Table 2 presents the results of the proposed model on the validation sets: each configuration is trained and evaluated 5 times to obtain the reported averages and standard deviations. Each experiment uses a different, randomly sampled and subject disjoint train-validation split and an SGD optimizer with learning rate of 10^{-3} and momentum of 0.9. We can immediately see that using random noise in this context totally misleads the neural network, and does not guide it towards a good choice of salient features. Secondly, we see that using Gaussian blur to preserve salient regions and "remove" non-iris regions is better than using nothing at all, both with iris normalization or cropping. This confirms our original intuition that the use of a low-pass filter on the patches covered by the contact lens could improve the matching scores compared to using the unfiltered images.

On the normalization versus cropping case, we can see that it is dependent on the saliency type used. When no saliency is applied, there was little difference whether the image was normalized or cropped. For images with noise in place of the textured contact lens, the cropped images had an advantage over normalized ones, yet considering the mediocre performance of the application of noise compared to the other runs, this improvement is insignificant. The more worthwhile comparison is between normalized and cropped versions with Gaussian blur: the experiments indicate a clear advantage to the runs with the normalized irises over the cropped ones, which is indeed an interesting finding, given that convolutional neural networks value spacial closeness and considering the normalization could be responsible for introducing undesired gradients in the input image.

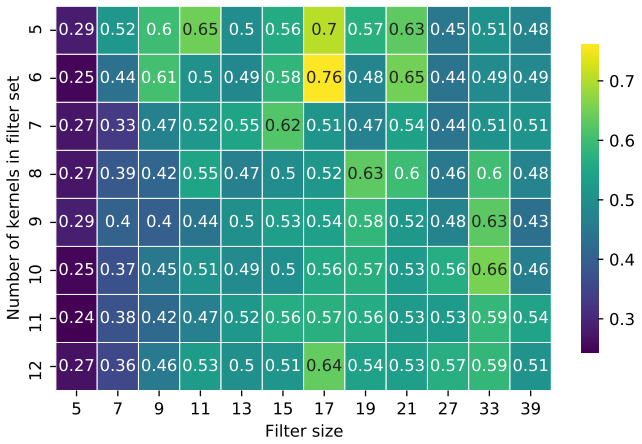

Figure 7: Heatmap of d' between genuine-impostor distributions as a function of the BSIF filter set: size and the number of kernels (translating to number of bits per pixel in BSIF encoding). The higher the value, the better the separation.

6. Comparison With State-Of-The-Art Iris Recognition

6.1. Human-Inspired BSIF-Based Iris Matcher

At least two recent papers [2, 14] suggested that human-inspired BSIF-based iris encoding [10] performs best among several open-source iris recognition methods today. We have thus decided to compare the proposed TCLA iris encoding and matching approach with the BSIF-based encoding and matching. It follow a dominant pipeline in iris recognition [12] (segmentation, normalization, filtering, binarization to get an iris code, and Hamming distance-based matching), however instead of Gabor wavelets, the ICA-based filters are applied that were derived from an experiment with humans solving the iris recognition task, and observed by an eye tracking device [10].

For a fair evaluation with TCLA matcher, which was designed with our training data, we select a new set of human-inspired BSIF filters, using *NDIris3D* (train + validation) dataset, to maximize the accuracy on this SOTA algorithm. Figure 7 shows d' obtained for all combinations of sizes and numbers of human-inspired BSIF filters. The highest d' was obtained for BSIF filter set composed of 6 filters of size 17×17. This configuration slightly outperformed the filter set proposed in the original paper (5 kernels of size 17×17). Thus, this new optimal filter set was used in experiments as the baseline counterpart for evaluating the performance of the TCLA matcher.

6.2. Evaluation On The Test Set

To compare the proposed TCLA matcher with the SOTA BSIF-based matcher, we performed the evaluation on the same testing partition of the *NDIris3D* dataset, whose samples remained unseen by the methods up to this stage. The recognition across the total of 1653 pairs was performed and their matching scores grouped into impostor and genuine distributions, over which were computed the d' (↑) and EER (↓).

Figure 1 shows the final comparison on the testing sets of the most suitable BSIF-based matcher (1-c) and the TCLA matcher using the best saliency guidance obtained on the validation set (1-d). The metrics indicate a strong advantage of our method over the traditional classifier, reducing the EER from 33.63% to 10.63% on the sequestered testing partition. This reduction in the error rate is confirmed by a d' of 2.10, still far from a separation under ideal conditions where it is common for matchers to reach $d' > 5$, yet it is 2.7 times larger than the BSIF-based matcher for the task at hand under the same circumstances.

7. Conclusions

In this paper, we presented a novel saliency-guided method based on Siamese Neural Networks that improves the recognition of subjects wearing textured contact lenses. The results on the testing set show a significant Equal Error Rate reduction by 23% when compared to the state-of-the-art human-inspired BSIF-based matcher, even after selecting the best performing BSIF filter set for that task. We have also presented and evaluated a model that is able to finely segment irises wearing textured contact lenses into their unaffected (authentic) and disguised (spoof) regions, an essential task for the creation of saliency maps used in the iris matcher.

Iris recognition with textured contact lenses remains a hard problem to solve. Samples with low to no authentic information cannot reliably be matched using the presented solution and might require investments on additional hardware for an improved data capture. Among the irises with enough information for usage, their quality might still present a great variance, which directly impacts: (i) how confident the matching can be, and (ii) what is the optimal feature extractor to use, given the properties of a sample..

The primary goal of this work was to propose an iris recognition method that maximizes the probability of a correct match (not increasing the false match rate at the same time) in case when subjects, correctly enrolled, are not conformant in a sense that they aim at disguising their identity by wearing texture contact lenses. While this research intentionally focuses on the usage of remaining iris patterns, a possible extension of this work is the usage

of ocular information to better discriminate subject pairs. This work may also support forensic human examiners, who may benefit from the segmentation model annotating alive parts of the iris, in case when the image of an iris wearing textured contact lens needs to be processed.

With this paper we also offer: (a) the TCLA segmentation model, (b) the TCLA matching model and (c) manual annotations of alive iris portions, corresponding to the texture contact lens samples taken from publicly-available *Notre Dame Photometric Stereo Iris Dataset*.

References

[1] Vijay Badrinarayanan, Alex Kendall, and Roberto Cipolla. SegNet: A Deep Convolutional Encoder-Decoder Architecture for Image Segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(12):2481–2495, 2017.

[2] Mauro Barni, Ruggero Donida Labati, Angelo Genovese, Vincenzo Piuri, and Fabio Scotti. Iris deidentification with high visual realism for privacy protection on websites and social networks. *IEEE Access*, 9:131995–132010, 2021.

[3] Amit Bendale, Aditya Nigam, Surya Prakash, and Phalguni Gupta. Iris Segmentation Using Improved Hough Transform. pages 408–415. 2012.

[4] Aidan Boyd, Kevin Bowyer, and Adam Czajka. Human-Aided Saliency Maps Improve Generalization of Deep Learning. 5 2021.

[5] Aidan Boyd, Kevin W. Bowyer, and Adam Czajka. Human-aided saliency maps improve generalization of deep learning. *CoRR*, abs/2105.03492, 2021.

[6] Aidan Boyd, Zhaoyuan Fang, Adam Czajka, and Kevin W. Bowyer. Iris presentation attack detection: Where are we now? *Pattern Recognition Letters*, 138:483–489, 2020.

[7] Canada Border Services Agency and U.S. Customs and Border Protection. NEXUS, 2018.

[8] Adam Czajka and Kevin W. Bowyer. Presentation attack detection for iris recognition: An assessment of the state of the art. *ACM Comput. Surv.*, 54(4):86:1–86:35, 2018.

[9] Adam Czajka, Zhaoyuan Fang, and Kevin W. Bowyer. Iris presentation attack detection based on photometric stereo features. *Proceedings - 2019 IEEE Winter Conference on Applications of Computer Vision, WACV 2019*, pages 877–885, 2019.

[10] A. Czajka, D. Moreira, K. Bowyer, and P. Flynn. Domain-specific human-inspired binarized statistical image features for iris recognition. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 959–967, Waikoloa Village, Hawai, United States, Jan 2019. IEEE.

[11] Priyanka Das, Joseph McFiratht, Zhaoyuan Fang, Aidan Boyd, Ganghee Jang, Amir Mohammadi, Sandip Purnapatra, David Yambay, Sebastien Marcel, Mateusz Trokielewicz, Piotr Maciejewicz, Kevin Bowyer, Adam Czajka, Stephanie Schuckers, Juan Tapia, Sebastian Gonzalez, Meiling Fang, Naser Damer, Fadi Boutros, Arian Kuijper, Renu Sharma, Cunjian Chen, and Arun Ross. Iris Liveness Detection Competition (LivDet-Iris) - The 2020 Edition. *IJCB 2020 - IEEE/IAPR International Joint Conference on Biometrics*, 2020.

[12] J. Daugman. How iris recognition works. *IEEE Trans. on Circuits and Systems for Video Tech.*, 14(1):21–30, January 2004.

[13] J. G. Daugman. High confidence visual recognition of persons by a test of statistical independence. *IEEE Trans. Pattern Anal. Mach. Intell.*, 15(11):1148–1161, November 1993.

[14] R. Donida Labati, Angelo Genovese, Vincenzo Piuri, Fabio Scotti, and Sarvesh Vishwakarma. I-social-db: A labeled database of images collected from websites and social media for iris recognition. *Image and Vision Computing*, 105:104058, 2021.

[15] Zhaoyuan Fang and Adam Czajka. Open Source Iris Recognition Hardware and Software with Presentation Attack Detection. In *IEEE International Joint Conference on Biometrics (IJCB)*, pages 1–8, Houston, Texas, Sept 28 – Oct 1, 2020.

[16] Z. Fang, A. Czajka, and K. W. Bowyer. Robust Iris Presentation Attack Detection Fusing 2D and 3D Information. *IEEE Transactions on Information Forensics and Security*, 16:510–520, 2021.

[17] Kaiming He, Georgia Gkioxari, Piotr Dollar, and Ross Girshick. Mask R-CNN. *Proceedings of the IEEE International Conference on Computer Vision*, 2017-Octob:2980–2988, 2017.

[18] Sheng Hsun Hsieh, Yung Hui Li, Wei Wang, and Chung Hao Tien. A novel anti-spoofing solution for iris recognition toward cosmetic contact lens attack using spectral ICA analysis. *Sensors (Switzerland)*, 18(3), 2018.

[19] Christian Rathgeb, Andreas Uhl, and Peter Wild. Incremental iris recognition: A single-algorithm serial fusion strategy to optimize time complexity. *IEEE 4th International Conference on Biometrics: Theory, Applications and Systems, BTAS 2010*, (819382), 2010.

[20] Yuxin Wu, Alexander Kirillov, Francisco Massa, Wan-Yen Lo, and Ross Girshick. Detectron2, 2019.

[21] Daksha Yadav, Naman Kohli, James S. Doyle, Richa Singh, Mayank Vatsa, and Kevin W. Bowyer. Unraveling the effect of textured contact lenses on iris recognition. *IEEE Transactions on Information Forensics and Security*, 9(5):851–862, 2014.