

Presentation Attack Detection for Cadaver Iris

Mateusz Trokielewicz

Institute of Control and Computation Engineering
Warsaw University of Technology
Nowowiejska 15/19, 00665 Warsaw, Poland

M.Trokielewicz@elka.pw.edu.pl

Adam Czajka

Biometrics and Machine Intelligence Lab
Research and Academic Computer Network
Kolska 12, 01045 Warsaw, Poland

adam.czajka@nask.pl

Piotr Maciejewicz

Department of Ophthalmology
Medical University of Warsaw
Lindleya 4, 02005 Warsaw, Poland

piotr.maciejewicz@wum.edu.pl

July 9, 2018

Abstract

This paper presents a deep-learning-based method for iris presentation attack detection (PAD) when iris images are obtained from deceased people. Post-mortem iris recognition, despite being a potentially useful method that could aid forensic identification, can also pose challenges when used inappropriately, i.e. utilizing a dead organ of a person in an unauthorized way. Our approach is based on the VGG-16 architecture fine-tuned with a database of 574 post-mortem, near-infrared iris images from the Warsaw-BioBase-PostMortem-Iris-v1 database, complemented by a dataset of 256 images of live irises, collected within the scope of this study. Experiments described in this paper show that our approach is able to correctly classify iris images as either representing a live or a dead eye in almost 99% of the trials, averaged over 20 subject-disjoint, train/test splits. We also show that the post-mortem iris detection accuracy increases as time since death elapses, and that we are able to construct a classification system with $APCER=0\%$ @ $BPCER\approx 1\%$ (Attack Presentation and Bona Fide Presentation Classification Error Rates, respectively) when only post-mortem samples collected at least 16 hours post-mortem are considered. Since acquisitions of ante- and post-mortem samples differ significantly, we applied countermeasures to minimize bias in our classification methodology caused by image properties that are not related to the PAD. This included using the same iris sensor in collection of ante- and post-mortem samples, and analysis of class activation maps to ensure that discriminant iris regions utilized by our classifier are related to properties of

the eye, and not to those of the acquisition protocol. This paper offers the first known to us PAD method in a post-mortem setting, together with an explanation of the decisions made by the convolutional neural network. Along with the paper we offer source codes, weights of the trained network, and a dataset of live iris images to facilitate reproducibility and further research.

1. Introduction

Post-mortem biometric identification is a field of study well established in the scientific community, among forensics professionals, but also in popular culture, with severed thumbs and plucked eyeballs being depicted in the big screen disturbingly often. With increasing importance that biometric authentication gains in our daily lives, fears are increasingly common among users, regarding the possibility of unauthorized access to our data, identity, or assets after our demise. Law enforcement officers in the U.S. are reportedly already using the fingerprints of the deceased to unlock the suspects' iPhones [32], which immediately brings up the topic of whether liveness detection should be one of the components of Presentation Attack Detection implemented in such devices. With a constantly growing market share of iris recognition, and recent research proving that iris biometrics in a post-mortem scenario can be viable [35, 34, 1, 18], these concerns are also becoming true for iris. In a recent interview for the *IEEE Spectrum* online magazine, Czajka discussed the issue of liveness detection, which is crucial in those cases when *we don't want* our biometric traits to be used after death [7].

To our best knowledge, there are no prior papers or pub-

lished research regarding the topic of discerning live irises from dead ones. This work thus offers the first study of iris liveness detection in a post-mortem scenario and offers the following **contributions to the state-of-the-art**:

- a method for iris liveness detection in a post-mortem setting, using a static iris image and based on a deep convolutional neural network (DCNN) fine-tuned with a dataset of post-mortem and live iris images,
- an analysis of iris regions or features being most used by the network when providing its decision, employing class activation mapping,
- a complementary dataset of live iris images, collected with the same equipment as the existing post-mortem iris images dataset,
- source codes and network weights for the offered solution.

Source codes, network weights, and a complementary dataset of live iris images can be obtained at <http://zbum.ia.pw.edu.pl/EN/node/46>.

This article is laid out as follows. Section 2 presents a short discussion of presentation attack detection methods for iris recognition, applications of deep learning for classification tasks, and techniques for explaining the reasoning of a DCNN-based classifier. Section 3 familiarizes the Reader with the dataset of live and post-mortem iris images used in this study. It also describes an initial experiment aiming at explaining the decisions that deep neural network is making when classifying iris images. Sections 4 and 5 provide an overview of the experimental methodology and results, respectively. Finally, relevant conclusions are given in Section 6.

2. Related work

2.1. Presentation Attack Detection in iris recognition

Presentation attack detection is already a well established area in the field of biometrics-related research. Existing methods include detection of fake representations of irises (paper printouts, textured contact lenses, prosthetic eyes, displays), or a non-conformant use of an actual eye. The most popular techniques used in iris PAD use various image texture descriptors (Binarized Statistical Image Features (BSIF) [16], Local Binary Patterns (LBP) [6], Binary Gabor Patterns (BGP) [17], Local Contrast-Phase Descriptor (LCPD) [11], Local Phase Quantization (LPQ) [28], Scale Invariant Descriptor (SID) [10], Scale Invariant Feature Transform (SIFT) and DAISY [21], Weber Local Descriptor (WLD) [11], or Wavelet Packet Transform (WPT) [2]), image quality descriptors [8], or deep-learning-based

techniques [19, 12, 21, 23]. If hardware adaptations are possible one may consider multi-spectral analysis [31] or estimation of three-dimensional iris features [20, 13] for PAD. Making the PAD more complex, one may consider measuring micro-movements of an eyeball, either using Eulerian video magnification [24] or by using an eye-tracking device [26], or measuring pupil dynamics [4]. An extensive review of the state of the art in PAD for iris recognition, including a systematization of attack methodologies and counter-measures proposed, can be found in [5], and independent evaluations of algorithms detecting iris paper printouts and textured contact lenses can be studied from the LivDet-Iris competition series (<http://livdet.org/>), which has had the last edition in 2017 [37].

Despite the abundance of research and proposed methods, there are still no published papers that would explore the concept of liveness detection in a scenario when cadaver (post-mortem) eyes are used to perform a presentation attack on the biometric sensor. However, one can still envisage such situation, in which a dead eye is used with an unsupervised biometric system to gain an unauthorized access to the assets the system is protecting.

2.2. Deep convolutional nets in image classification

Deep convolutional neural networks (DCNN) are useful in solving selected groups of computer vision problems, such as image classification [29], semantic segmentation [9], automatic image captioning [15], or visual question answering [25].

These methods learn model parameters (network weights and biases) and hyper-parameters (network architecture details and training constraints) by guessing them from the data that is made available to the model during the training phase. Thus, the model learns the data itself, constituting a *feature-learned* or *data-driven* approach. This is opposed to finding the model parameters by obtaining prior knowledge about the object being recognized and fine-tuning much fewer parameters of a less complicated model, which may be called *feature engineering* or *hand-crafted* approaches. Data-driven approaches, such as these involving neural networks, offer significant advantages over hand-crafted ones in situations where the knowledge about the subject is either limited or difficult to be put into relatively simple mathematical rules that would enable building a feature-engineered solver.

This can be the case with processing post-mortem iris samples, either to recognize a person or to recognize a presentation attack. Trokielewicz *et al.* revealed that although post-mortem iris recognition is, to some extent, possible with current software, it also poses new challenges that we do not yet have solutions to [34]. Most importantly, there currently are no mathematical models that would explain the iris' behavior over the course of post-mortem time hori-

zon, *i.e.* quantify and predict the changes that the iris may undergo after one’s demise. Therefore, when aiming at discerning live irises from dead ones, a potentially promising way of solving this problem is to rely on the feature-learned approach that utilizes the existing datasets of post-mortem and live iris images to teach itself to give the correct answer.

This has already proven promising by enabling us to propose a post-mortem iris image segmentation method based on a DCNN, which achieves performance superior to the typical iris recognition method [33].

2.3. Explaining deep networks’ decisions

Despite spectacular successes in computer vision tasks, DCNNs all have a significant drawback, namely their inability to provide an intuitive, human-understandable explanation for their decisions. Hence, although the performance of these solutions may be excellent, in their basic designs they lack interpretability. This has at least two potential downsides: 1) if humans do not know how an expert system (*e.g.* a self-driving car software) works, they will not trust it, and 2) if the creators do not know how the system works, they cannot improve it.

Zhou *et al.* [38] propose a technique called *class activation mapping*, or CAM, for identification of discriminative image regions, *i.e.* these, which are decisive when it comes to the classification output. This is said to be feasible despite training the network only with image-level labels, and not dense, pixel-wise labels. The authors achieve this by dropping fully-connected layers in the popular network architectures (ALexNet, GoogLeNet, VGG), and replacing them with global average pooling layers followed only by a fully-connected output softmax layer. To increase spatial resolution of the mappings, some convolutional layers are also removed from the architectures, resulting in 14×14 or 15×15 outputs from the convolutional part of the respective network. This approach enables highlighting image regions that are important for discrimination, and even localize regions responsible for detecting patterns, such as text, and even higher-order concepts.

Selvaraju *et al.* introduce improvement over the Zhou’s method with Grad-CAM [27], a technique similar to CAM, but not requiring any changes to the network’s architecture, and thus easier to use on the already trained models. This approach produces coarse localization heat-maps highlighting the regions that are considered important by the network when predicting the concept in the image. Also, by combining these low-resolution maps with high-resolution visualization of features learned by the network, obtained from guided back-propagation introduced by Springenberg *et al.* in [30], it is possible to obtain a more fine-grained importance maps, which apart from highlighting a coarse *region* of the image that is considered discriminatory, also allows insight into which *features* are important.

These techniques can be important for two reasons, both of them being explored in our paper. First, class activation mapping can help analyze the potential bias in the raw data, that can interfere with the network training, causing the model to learn features that are not directly related to the task at hand. For instance, learning the presence of metal retractors used to open cadaver eyes, and missing in live eyes, ends up with perfect accuracy albeit with no relation to PAD accuracy. Second, we hope to gain some knowledge regarding the iris/eye features being employed by the network for discriminating between live and dead irises.

3. Experimental dataset

3.1. Post-mortem iris images

For the purpose of this study, we used the only, known to us, publicly available Warsaw-BioBase-PostMortem-Iris-v1 dataset, which gathers 1,330 post-mortem iris images collected from 17 individuals during various times after death (from 5 hours up to 34 days) [36]. These samples represent ocular regions of recently deceased subjects. In addition to typical, near-infrared (NIR) iris images collected from cadavers, high quality visible light images are also available, however, for the purpose of this study we only employed the former, as NIR samples are usually employed in commercial, deployed iris recognition systems. There are 574 NIR images available in the dataset, most of them captured up to two days after death, but some of the samples extend up to 814 hours after death.

3.2. Images of live irises

Since post-mortem part of the dataset used in this study does not offer any ante-mortem samples, or reference images of live individuals, we had to collect a complementary dataset of iris images collected from live people. To mimic the original acquisition protocol as closely as possible, and thus to minimize the bias in training the DCNN, we have employed the same iris camera as was used in the post-mortem counterpart, namely the IriTech IriShield M2120U.

3.3. Analysis of potential data bias

We are aware that even despite our best efforts to keep the acquisition protocols as alike as possible, there will be some bias in the data. The aim of this Section is thus to carefully examine these differences, discuss their importance and impact on the experiments, and to propose countermeasures, where possible. To quantify the variations between images coming from the two datasets that may originate in difference between camera operators or the environment, we have performed the calculations of three covariates related to iris quality, namely: average intensity, grayscale utilization (image histogram entropy), and image sharpness, the latter two being suggested by the ISO/IEC standard on

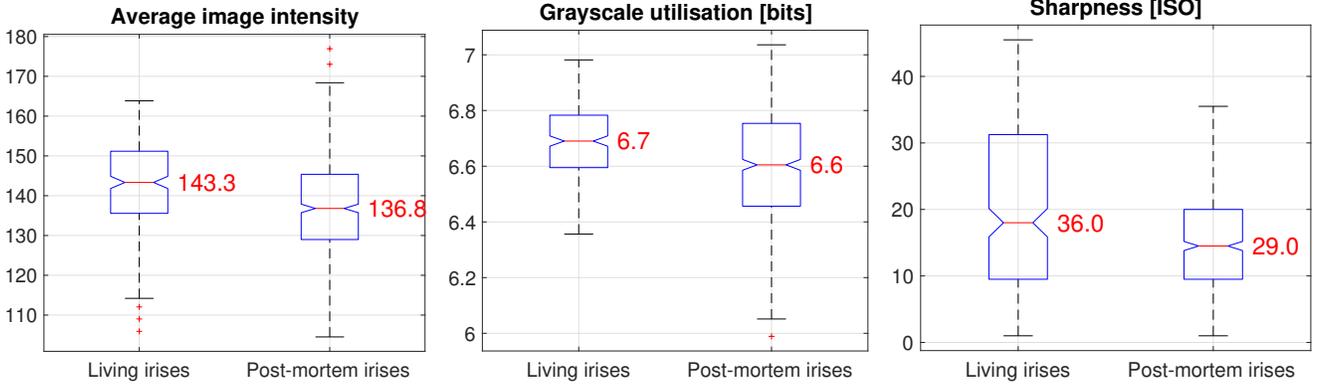


Figure 1: Boxplots representing differences in image quality metrics between the two datasets. Median values are shown in red, height of each boxes corresponds to an inter-quartile range (IQR) spanning from the first (Q1) to the third (Q3) quartile, whiskers span from $Q1-1.5*IQR$ to $Q3+1.5*IQR$, and outliers are shown as crosses.

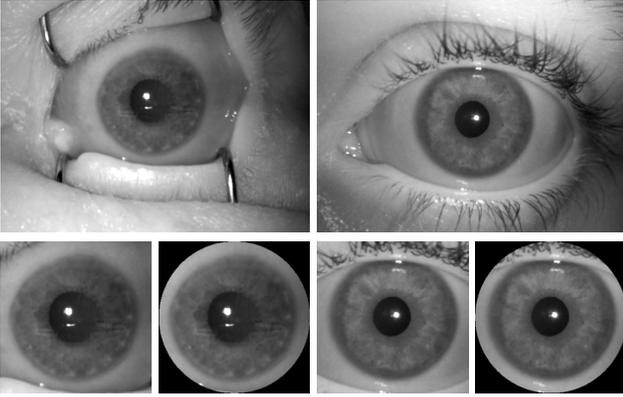


Figure 2: Example iris images obtained from a dead (**left**) and a live subject (**right**): original images (**top row**) and their cropped and cropped_masked versions (**bottom row**).

iris image quality [14]. These covariates are defined as follows:

- **average image intensity:**

$$AI = \frac{1}{N} \sum_{i=1}^n \sum_{j=1}^m I_{ij}$$

where I_{ij} is the pixel intensity and n, m is the image size in pixels;

- **grayscale utilization**, or the entropy of the iris image histogram H measured in bits, examines pixel values of an iris image to calculate a spread of intensity values and assess whether the image is properly exposed:

$$H = - \sum_{i=1}^{256} p_i \log_2 p_i$$

where p_i is the probability of each gray level i occurring in the image, hence, the total count of pixels at gray level i , divided by the total number of pixels in the image [14];

- **sharpness**, determined by the power resulting from filtering the image with a Laplacian of Gaussian kernel; for brevity, we do not reproduce the formulae here, and instead refer the Reader to the original ISO/IEC documentation [14].

Results of these calculations are shown in Fig. 1. Notably, only the *sharpness* covariate differs largely between the two datasets, as post-mortem iris images have lower sharpness on average. This can be a result of a combination of factors, such as: more difficult collection environment (*e.g.* a hospital mortuary), a less experienced operator (*e.g.* medical staff), limited time, and such. For completeness, to provide formal statistical analysis, we ran a Wilcoxon rank-sum test for each pair of covariates, which revealed that there are statistically significant differences between the subsets of live and post-mortem iris images, as the null hypothesis stating that the compared scores are samples from continuous distributions with equal median was rejected at significance level $\alpha = 0.05$ in all three cases. However, all three covariates do not provide enough differentiation between the images coming from different datasets, to themselves serve as features for presentation attack detection.

3.4. Assessing data bias with Grad-CAM

Despite our efforts to follow the same collection guidelines as those of [36], there are still some differences between the respective samples that originate in different presentation of the biometric itself in live and post-mortem scenarios. This is most notably related to the appearance

of eyelids, which in the post-mortem data are often pulled apart with a metal retracting device to keep the eye open for image acquisition. To at least partially mitigate these differences, the subjects participating in the collection of the reference data were asked to open their eyes as widely as possible. However, the presence of metal parts of the medical equipment is still an issue, as these appear in post-mortem cases and do not appear in live cases.

To examine whether these metal retractors will serve as cues for the DCNN when it is trained to discern post-mortem irises from the live ones, we have employed the class activation mapping technique described in Section 2.3 by performing an additional experiment, in which the network, same as described later in Sec. 4, was trained with the original, uncropped images, as shown in Fig. 2, top row. This initial implementation of the network was done in Keras [3], using an adapted code from [22].

Example predictions were then obtained for both the *live* and *post-mortem* classes, Figs. 3 and 5. This shows that the metal parts of the medical equipment used to open the eyelids of deceased subjects indeed **provide class discriminatory cues** (cf. top row in Fig. 3), which in this case is strongly undesirable, as we want our network to recognize post-mortem irises, and not the equipment that may, or may not accompany the post-mortem data collection. Also, for the second investigated post-mortem sample, for which no metal parts are visible, but instead a heavily distorted eyelids are present, the network also pays attention not to the iris itself, but rather to its surrounding – in this case, the eyelids, cf. bottom row in Fig. 3. In none of these cases can we see strong activations by the iris region. However, when distinctive features such as metal retractors or heavily distorted eyelids are absent from the image, and its overall resemblance to an image of a live iris is strong, the model fails to make a correct prediction, and focuses on the iris region instead, Fig. 4.

Notably, a similar behavior can be observed for images of live irises. When analyzing example activation maps for the *live* class, we see that it is the iris region that produces the strongest activations, while the iris surroundings remain mostly unused, as depicted in Fig. 5. This behavior, contrary to what we observe with post-mortem samples, can be considered desirable.

3.5. Dataset modification to counteract the bias

To force-shift the network’s attention to the iris, and not its neighborhood, we have manually segmented all images in both datasets, approximating the outer iris boundary with a circle with a radius R_i and then cropping and masking the image to the size of $1.2R_i$, see Fig. 2. This margin of $0.2R_i$ is preserved purposefully, to represent the differences in the iris surroundings and in the iris-sclera boundary, and not only these present in the tissue itself. The same rea-

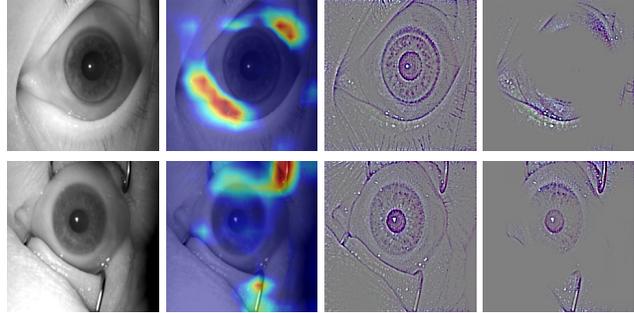


Figure 3: Example class activation maps obtained using the Grad-CAM technique for samples from the original, unmodified dataset, with a model trained on the original dataset. From left to right: (1) original image, (2) Grad-CAM, (3) guided back-propagation, (4) a combination of (2) and (3). **Post-mortem samples are represented here, both correctly classified.**

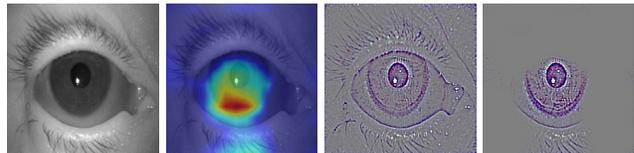


Figure 4: Same as in Fig. 3, but this **post-mortem sample was misclassified as a live one.**

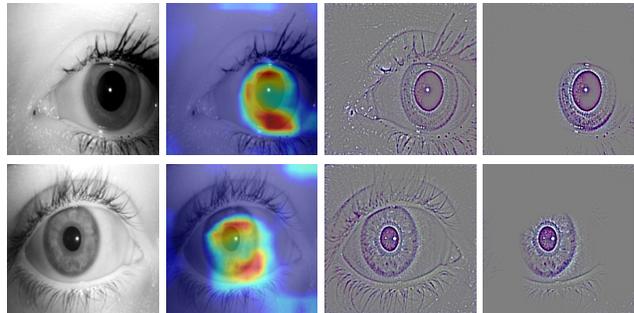


Figure 5: Same as in Fig. 3, but for **samples of live irises.**

soning is behind leaving the pupillary region unmasked, as the appearance of the pupil can bear liveness information as well. The raw images obtained from the sensor are thus referred to as the *original dataset*, while the modified version is called the *cropped_masked dataset*.

To validate this reasoning, we train the same model that was employed for assessing class activation maps in the unmodified data, but with the *cropped_masked* images instead. Activation maps drawn for example *cropped_masked* iris images are shown in Figs. 6, 7, and 8. As for the *post-mortem* samples, the new model now mostly focuses on the iris and its boundary, cf. Fig. 6, which seems reasonable, as the iris-sclera boundary is quickly getting blurry as time

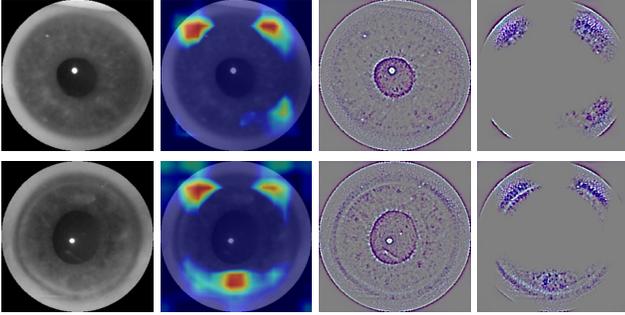


Figure 6: Example class activation maps obtained using the Grad-CAM technique for samples from the *cropped_masked* dataset, with a model trained on the *cropped_masked* dataset as well. From left to right: (1) original image, (2) Grad-CAM, (3) guided back-propagation, (4) a combination of (2) and (3). **Post-mortem samples are represented here, both correctly classified.**

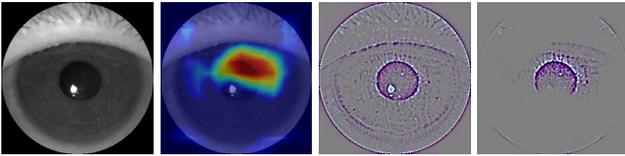


Figure 7: Cropped version of the post-mortem sample misclassified earlier in Fig. 4. The classification is now **correct**.

since death progresses. This is different for the problematic sample examined earlier in Fig. 4, for which the activation map is centered near the pupillary region of the eye. This sample, however, is now correctly classified as a post-mortem one.

When it comes to samples representing live eyes, the network also seems to assess their PAD score by analyzing the iris-boundary region, but to some degree it also brings its attention to the iris itself, and to the specular reflection found in the middle of the pupil, as depicted in Fig. 8. These features seem to offer enough discriminatory power to successfully differentiate between the post-mortem and live samples, as the validation accuracy on the *cropped_masked* subset of subject-disjoint iris images reaches 100%, compared to less than 95% obtained for the unmodified set of the same iris images. This also shows that we have managed to successfully shift the attention of the network towards the iris features and discriminative information they offer in recognizing cadaver samples.

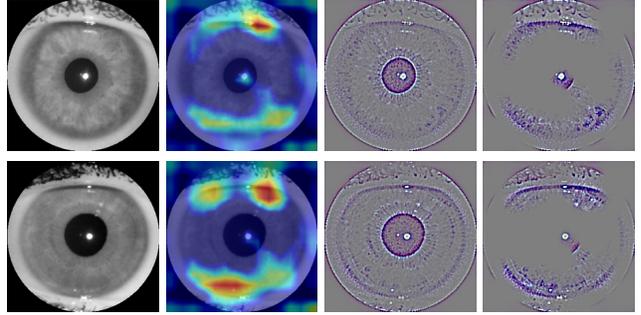


Figure 8: Same as in Fig. 6, but for **samples of live irises**.

4. Proposed methodology and evaluation

4.1. Model architecture

For our solution, we employed the VGG-16 model pre-trained on natural images from the ImageNet database [29], which has been shown to repeatedly achieve excellent results in various classification tasks after minor adaptation and re-training. We thus performed a simple modification to the last three layers of the original graph to reflect the nature of our binary classification into *live* and *post-mortem* types of images, and performed transfer learning by fine-tuning the network weights to our dataset of iris images representing both classes.

4.2. Training and evaluation procedure

For the network training and testing procedure, 20 subject-disjoint train/test data splits were created by randomly assigning the data from 3 subjects to the test subset, and the data from the remaining subjects to the train subset, both for the *live* and *post-mortem* parts of the database. These twenty splits were made with replacement, making them statistically independent. The network was then trained with each train subset independently for each split, and evaluated on the corresponding test subset. This procedure gives 20 statistically independent evaluations and allows to assess the variance of the estimated error rates. The training, encompassing 10 epochs in each of the train/test split run, was performed with stochastic gradient descent as the minimization method with momentum $m = 0.9$ and learning rate of 0.0001, with the data being passed through the network in mini batches of 16 images.

During testing, a prediction of the *live* or *post-mortem* class-wise probability was obtained from the Softmax layer, together with a corresponding predicted categorical label. The probabilities of *post-mortem* samples belonging to their class are also associated with a time that elapsed since death until the sample acquisition. This allows the analysis of classification accuracy in respect to the *post-mortem* time horizon.

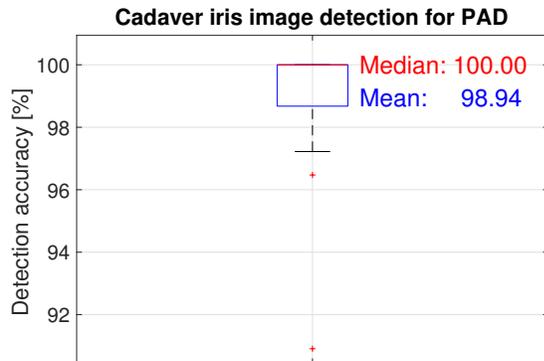


Figure 9: Accuracy of classification into live and post-mortem classes, achieved in 20 independent, subject-disjoint train/test data splits.

5. Results

5.1. Averaged classification accuracy

Several metrics can be utilized to evaluate the classification accuracy of our solution. First, we average the accuracy achieved in each of the twenty train/test splits, measured as a share of correctly assigned labels to the overall number of trials in a given split, Fig. 9. Notably, in most of the splits, the solution achieves a 100% classification accuracy on the test subset, with the average of 98.94%.

5.2. Global performance by Receiver Operating Characteristic

Receiver Operating Characteristic curves (ROC) are often employed for visualizing performance of classification systems, especially when binary classification is in place. Thus, we present ROC graphs with Areas Under Curve (AUC) calculated for each curve in Fig. 12. The obtained classification accuracy is almost perfect on the whole set of samples (AUC=99.94%), and even better when samples acquired shortly (*i.e.* five hours) after death are excluded from the dataset (AUC=99.99%). We elaborate more on this exclusion in the following paragraph.

5.3. Classification accuracy in respect to post-mortem time horizon

Having post-mortem samples acquired for different subjects at multiple time points after death, makes it possible to analyze the performance of our solution in respect to the time that has passed since a subject's demise, Fig. 10. Interestingly, and perhaps accordingly to a common sense, the probability of a sample being post-mortem increases as time since death elapses. We can expect a few false matches (post-mortem samples being classified as live iris samples) with images obtained 5 hours after death, regardless of the chosen threshold. This can be attributed to

the fact that these images are very similar to those obtained from live individuals, as post-mortem changes to the eye are still not pronounced enough to allow for a perfect classification accuracy. Score means and standard deviations in a close-up of the time period occurring shortly after death is shown in Fig. 11, which shows that the increased variations in the score distributions in several time points, with the largest variance occurring 5 hours after death. However, the already good accuracy is getting close-to-perfect when these samples are not taken into consideration, Fig. 9, red dotted line. Attack Presentation Classification Error Rate (APCER, misclassifying post-mortem samples as live ones) equal to zero can be achieved in such scenario, provided that an appropriate acceptance threshold is established. Bona Fide Presentation Classification Error Rate (BPCER, misclassifying live samples as post-mortem ones) in such case is only about 1%.

6. Conclusions

This paper offers the first, know to us, method for iris liveness detection in respect to the post-mortem setting, based on a deep convolutional neural network VGG-16, adapted and fine-tuned to the task of discerning live and dead irises. The proposed method is able to correctly classify nearly 99% of the samples, assigning *alive* or *post-mortem* labels, respectively.

Another interesting insight is that samples collected briefly after death (*i.e.* five hours in our study) can fail to provide post-mortem changes that are pronounced enough to serve as cues for liveness detection. However, when classifying samples collected at least 16 hours after death, we found that employing a well-tuned threshold enables APCER=0%, meaning that no post-mortem sample gets mistakenly classified as a live one, with a probability of misclassifying a live sample as a dead one being around BPCER=1%. This shows that while post-mortem iris images are relatively easy to identify, those obtained very shortly after a subject's demise can pose problems for automatic solutions due to post-mortem changes not being prominent enough yet.

A significant portion of this paper is also dedicated to an attempt to explain the reasoning behind the decisions provided by our DCNN-based solution, especially with respect to the iris regions that are considered when making these decisions. By employing the Grad-CAM class activation mapping technique, we managed to show that the image regions bearing the most useful discriminatory cues for liveness detection are those containing iris-sclera boundary, and to some extent also the pupillary region.

We follow the guidelines on research reproducibility by providing a) the source codes, b) trained network weights, and c) the complementary dataset of live irises to all interested researchers, who would like to reproduce the results

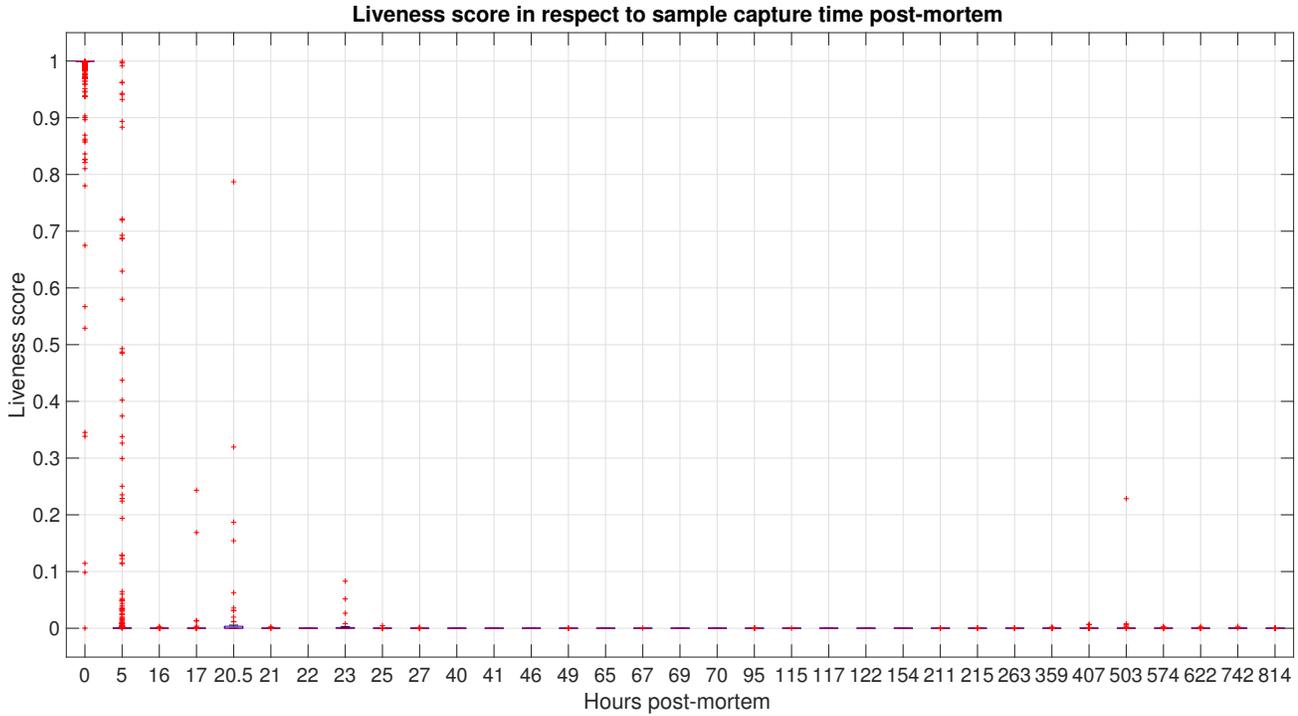


Figure 10: Boxplots representing differences in liveness scores earned by the samples acquired in different moments after death. **The lower the score, the more likely the sample represents a post-mortem eye.** Samples denoted as acquired **zero hours** post-mortem are those collected from live irises.

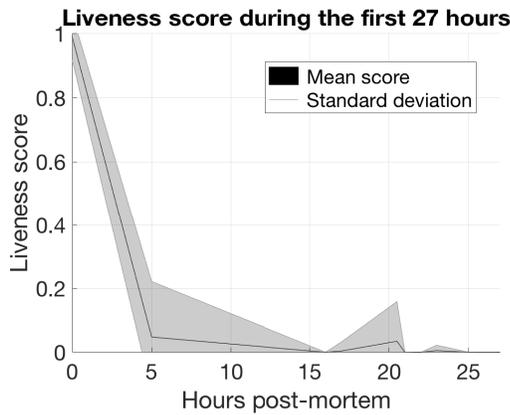


Figure 11: Means (solid black line) and standard deviations (grey areas) of liveness score predicted by the network for samples obtained up to 27 hours after death. Samples denoted as acquired **zero hours** post-mortem are those collected from live irises.

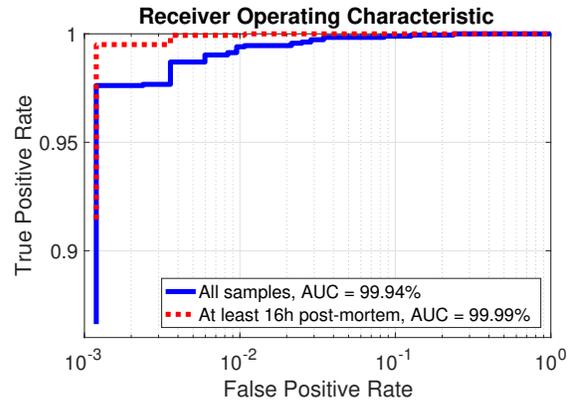


Figure 12: ROC curve showing classification accuracy averaged over 20 independent, subject-disjoint train/test data splits. AUC denotes the Area Under Curve metric.

Acknowledgment

presented in this paper. We hope that this contribution will stimulate further research in iris presentation attack detection.

The authors would like to thank Kasia Roszczewska for allowing us to execute the calculations on her GPU, as well as NVIDIA for providing a GPU unit for our laboratory.

References

- [1] D. S. Bolme, R. A. Tokola, C. B. Boehnen, T. B. Saul, K. A. Sauerwein, and D. W. Steadman. Impact of environmental factors on biometric matching during human decomposition. In *2016 IEEE 8th International Conference on Biometrics Theory, Applications and Systems (BTAS)*, pages 1–8, Sept 2016.
- [2] R. Chen, X. Lin, and T. Ding. Liveness detection for iris recognition using multispectral images. *Pattern Recognition Letters*, 33(12):1513 – 1519, 2012.
- [3] F. Chollet. Keras: Deep Learning library for Theano and TensorFlow. <https://keras.io/>, 2015.
- [4] A. Czajka. Pupil dynamics for iris liveness detection. *IEEE Trans. Inf. Forens. Security*, 10(4):726–735, April 2015.
- [5] A. Czajka and K. Bowyer. Presentation Attack Detection for Iris Recognition: An Assessment of the State of the Art. *ACM Computing Surveys (in review)*, <https://arxiv.org/abs/1804.00194>, 2018.
- [6] J. S. Doyle, P. J. Flynn, and K. W. Bowyer. Automated classification of contact lens type in iris images. In *IEEE Int. Conference on Biometrics (ICB)*, pages 1–6, June 2013.
- [7] Eliza Strickland. Biometrics Researcher Asks: Is That Eyeball Dead or Alive?, <https://spectrum.ieee.org/the-human-os/biomedical/imaging/biometric-researcher-asks-is-that-eyeball-alive-or-dead> (accessed: March 20, 2018).
- [8] J. Galbally, M. Savvides, S. Venugopalan, and A. A. Ross. *Iris Image Reconstruction from Binary Templates*, pages 469–496. Springer London, London, 2016.
- [9] A. Garcia-Garcia, S. Orts-Escolano, S. Oprea, V. Villena-Martinez, , and J. Garcia-Rodriguez. A Review on Deep Learning Techniques Applied to Semantic Segmentation. <https://arxiv.org/abs/1704.06857v1>, 2017.
- [10] D. Gragnaniello, G. Poggi, C. Sansone, and L. Verdoliva. Contact lens detection and classification in iris images through scale invariant descriptor. In *Int. Conference on Signal-Image Technology Internet-Based Systems (SITIS)*, pages 560–565, Nov 2014.
- [11] D. Gragnaniello, G. Poggi, C. Sansone, and L. Verdoliva. An investigation of local descriptors for biometric spoofing detection. *IEEE Trans. Inf. Forens. Security*, 10(4):849–863, Apr. 2015.
- [12] L. He, H. Li, F. Liu, N. Liu, Z. Sun, and Z. He. Multi-patch convolution neural network for iris liveness detection. In *IEEE Intl. Conference on Biometrics Theory, Applications and Systems (BTAS)*, pages 1–7, September 2016.
- [13] K. Hughes and K. W. Bowyer. Detection of contact-lens-based iris biometric spoofs using stereo imaging. In *Hawaii Intl. Conference on System Sciences*, pages 1763–1772, January 2013.
- [14] ISO/IEC 29794-6. Information technology – Biometric sample quality - Part 6: Iris image data (FDIS), September 2014.
- [15] J. Johnson, A. Karpathy, and L. Fei-Fei. DenseCap: Fully Convolutional Localization Networks for Dense Captioning. *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [16] J. Komulainen, A. Hadid, and M. Pietikinen. Generalized textured contact lens detection by extracting bsif description from cartesian iris images. In *IEEE Int. Joint Conference on Biometrics (IJCB)*, pages 1–7, Sept 2014.
- [17] Lovish, A. Nigam, B. Kumar, and P. Gupta. Robust contact lens detection using local phase quantization and binary gabor pattern. In G. Azzopardi and N. Petkov, editors, *Int. Conference on Computer Analysis of Images and Patterns (CAIP)*, pages 702–714. Springer International Publishing, 2015.
- [18] P. M. Mateusz Trokielewicz, Adam Czajka. Iris Recognition After Death. *IEEE Transactions on Information Forensics and Security (in review)*, <https://arxiv.org/abs/1804.01962>, 2018.
- [19] D. Menotti, G. Chiachia, A. Pinto, W. R. Schwartz, H. Pedrini, A. X. Falcão, and A. Rocha. Deep representations for iris, face, and fingerprint spoofing detection. *IEEE Transactions on Information Forensics and Security*, 10(4):864–879, April 2015.
- [20] A. Pacut and A. Czajka. Aliveness detection for iris biometrics. In *IEEE Int. Carnahan Conference on Security Technology (ICCST)*, pages 122–129, October 2006.
- [21] F. Pala and B. Bhanu. Iris liveness detection by relative distance comparisons. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, July 2017.
- [22] V. Petsiuk. Keras implementation of GradCAM. <https://github.com/eclique/keras-gradcam>, accessed on April 4, 2018.
- [23] R. Raghavendra, K. B. Raja, and C. Busch. Contlensnet: Robust iris contact lens detection using deep convolutional neural networks. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1160–1167, March 2017.
- [24] K. Raja, R. Raghavendra, and C. Busch. Video presentation attack detection in visible spectrum iris recognition using magnified phase information. *IEEE Trans. Inf. Forens. Security*, 10(10):2048–2056, October 2015.
- [25] M. Ren, R. Kiros, and R. S. Zemel. Exploring models and data for image question answering. *CoRR*, abs/1505.02074, 2015.
- [26] I. Rigas and O. V. Komogortsev. Eye movement-driven defense against iris print-attacks. *Pattern Recognition Letters*, 68, Part 2:316 – 326, 2015. Special Issue on Soft Biometrics.
- [27] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra. Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization. *ICCV*, <https://arxiv.org/abs/1610.02391>, 2016.
- [28] A. F. Sequeira, S. Thavalengal, J. Ferryman, P. Corcoran, and J. S. Cardoso. A realistic evaluation of iris presentation attack detection. In *Int. Conference on Telecommunications and Signal Processing (TSP)*, pages 660–664, June 2016.
- [29] K. Simonyan and A. Zisserman. Very Deep Convolutional Networks for Large-Scale Image Recognition. <https://arxiv.org/abs/1409.1556>, 2014.
- [30] J. T. Springenberg, A. Dosovitskiy, T. Brox, and M. Riedmiller. Striving for Simplicity: The All Convolutional Net. *ICLR*, <https://arxiv.org/abs/1412.6806>, 2015.
- [31] S. Thavalengal, T. Nedelcu, P. Bigioi, and P. Corcoran. Iris liveness detection for next generation smartphones. *IEEE Trans. Consumer Electronics*, 62(2):95–102, May 2016.

- [32] Thomas Brewster. Yes, Cops Are Now Opening iPhones With Dead People's Fingerprints, <https://www.forbes.com/sites/thomasbrewster/2018/03/22/yes-cops-are-now-opening-iphones-with-dead-peoples-fingerprints> (accessed: March 20, 2018).
- [33] M. Trokielewicz and A. Czajka. Data-driven Segmentation of Post-mortem Iris Images, 6th IAPR/IEEE International Workshop on Biometrics and Forensics (IWF 2018), June 7-8, 2018, Sassari, Italy, 2018.
- [34] M. Trokielewicz, A. Czajka, and P. Maciejewicz. Human Iris Recognition in Post-mortem Subjects: Study and Database, 8th IEEE International Conference on Biometrics: Theory, Applications and Systems (BTAS 2016), September 6-9, 2016, Buffalo, NY, USA, 2016.
- [35] M. Trokielewicz, A. Czajka, and P. Maciejewicz. Post-mortem Human Iris Recognition, 9th IAPR International Conference on Biometrics (ICB 2016), June 13-16, 2016, Halmstad, Sweden, 2016.
- [36] Warsaw University of Technology. Warsaw-BioBase-PostMortem-Iris-v1.0: <http://zbum.ia.pw.edu.pl/en/node/46>, 2016.
- [37] D. A. Yambay, B. Becker, N. Kohli, D. Yadav, A. Czajka, K. W. Bowyer, S. Schuckers, R. Singh, M. Vatsa, A. Noore, D. Gragnaniello, C. Sansone, L. Verdoliva, L. He, Y. Ru, H. Li, N. Liu, Z. Sun, and T. Tan. LivDet 2017 - Iris Liveness Detection Competition 2017. *Biometrics: Theory Applications and Systems (BTAS)*, pages 0–5, 2017.
- [38] B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba. Learning Deep Features for Discriminative Localization. *CVPR*, <https://arxiv.org/abs/1512.04150>, 2016.