



Impact of Blinking on Deep Learning Based Iris Recognition

Daniel Tebor

Kennesaw State University

Marietta, Georgia, USA

dtebor@students.kennesaw.edu

Eli Headley

Kennesaw State University

Marietta, Georgia, USA

eheadle1@students.kennesaw.edu

Mahmut Karakaya

Kennesaw State University

Marietta, Georgia, USA

mkarakay@kennesaw.edu

ABSTRACT

This study explores the impact of blinking on deep learning based iris recognition, addressing a critical aspect in the development of robust, reliable, and non-intrusive biometric systems. While previous research has demonstrated the promise of Convolutional Neural Networks (CNNs), such as AlexNet, GoogleLeNet, and ResNet, the impact of blinking remains underexplored in this context. To address this gap, our research focuses on training multiple ResNet models with varying degrees of iris occlusion exposure. Using a dataset with 101 subjects, we generated cohorts of synthetically occluded images ranging from 0% occlusion to 90% occlusion. Our findings reveal a noteworthy linear performance decrease in models unexposed to blinked images as iris occlusion increases. However, augmenting the training dataset with occluded images significantly mitigates this performance degradation, highlighting the importance of accounting for blinking in the development of reliable iris recognition systems.

CCS CONCEPTS

• Computing methodologies → Object identification; Image segmentation; Interest point and salient region detections.

KEYWORDS

Biometrics, Iris Recognition, Eye Blink, Deep Learning, Convolutional Neural Networks

ACM Reference Format:

Daniel Tebor, Eli Headley, and Mahmut Karakaya. 2024. Impact of Blinking on Deep Learning Based Iris Recognition. In *2024 ACM Southeast Conference (ACMSE 2024), April 18–20, 2024, Marietta, GA, USA*. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3603287.3651206>

1 INTRODUCTION AND BACKGROUND

Biometric-driven identification has become an increasingly valuable tool in modern applications. Among different biometric systems, iris recognition is valuable as irises do not significantly change with age or emotional state. Furthermore, iris recognition is discrete, accurate, and non-intrusive even when part of a subject's face is obscured or contorted. However, despite these advantages, traditional iris recognition methods face challenges, primarily stemming from the limitations of algorithms employed in earlier research [3].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ACMSE 2024, April 18–20, 2024, Marietta, GA, USA

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0237-2/24/04...\$15.00

<https://doi.org/10.1145/3603287.3651206>

The initial approaches relied on traditional algorithms, restricting reliable usage to only frontal images captured under specific conditions due to variations in iris geometry with viewing angles. Despite attempts to rectify distortions through techniques like segmentation using active contours [1] and perspective transformations [9], achieving consistent and reliable results remained challenging.

The limitations of traditional iris recognition methods led to an exploration of convolutional neural networks (CNNs) for iris biometric systems. The previous investigations of CNNs included models such as AlexNet [6], GoogleLeNet [10], and ResNet50 [4] using transfer learning [5]. These studies revealed that training with off-angle images and incorporating features around the iris led to improved performance, particularly for images captured at non-frontal angles. Since most images in these datasets had limited iris images with occlusion caused by blinking, they ignored the effect of blinking on deep learning based iris recognition.

Blinking emerges as a critical factor in deep learning based iris recognition, especially in establishing robust standoff biometric systems. A minimally intrusive system that accommodates natural blinking while ensuring accuracy is imperative for enabling subjects to move freely at a distance. While excluding highly occluded images, such as those captured during complete blinking, may seem feasible, subjects are unlikely to keep their eyes wide open without specific instructions. Coupled with other factors like gaze angle, this could lead to suboptimal image conditions for standoff iris recognition in a majority of real-world situations. Therefore, to create standoff iris biometric systems that are reliable under a variety of conditions, it is important to investigate the effects of blinking on iris recognition and to explore how performance can be improved in situations where iris occlusion is not optimal.

In literature, the existing research on the impact of blinking in deep learning iris recognition remains relatively limited. Contemporary studies, employing both traditional and deep learning methods, have predominantly focused on investigating factors other than such as pupil dilation. However, there has been little research related to blinking. Liu et al., for instance, proposed a methodology involving sequential forward floating selection (SFFS) for optimizing feature selection and a non-linear support vector machine (SVM) for identification [7]. Their findings present promising results, demonstrating efficacy for both blinked and unblinked images. Our research varies from this approach, instead using out-of-box models modified using transfer learning as the foundation for iris recognition, and then elaborating on the baseline by augmenting using blinked images.

This paper extends the exploration of CNNs with transfer learning in the context of iris recognition. Our research contributions include i) investigating the impact of blinking on iris recognition using pre-trained CNN models, ii) augmenting CNN training with synthetically occluded iris images to enhance performance, and iii)

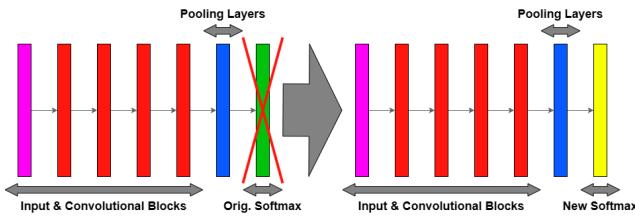


Figure 1: Flowchart of Transfer Learning with ResNet50

quantifying the tradeoffs associated with training CNNs using such synthetic occlusions. Our research contributes valuable insights to guide the evolution of biometric iris recognition for more versatile and robust implementations.

2 METHODOLOGY

Our methodology includes three main denominations: i) the selection and modification of a deep learning model through transfer learning, ii) the modification of our existing dataset, and iii) designing experiments for the testing of iris occlusion.

2.1 Deep Learning Framework

The research utilized a Convolutional Neural Network (CNN), specifically ResNet50 [4], for feature extraction from iris images and classification of different irises. ResNet50 was selected due to its demonstrated effectiveness [5] when compared to other models such as AlexNet [6] and GoogleLeNet [10]. The model's input is a 3-dimensional tensor with a shape of 3 by 224 by 224, representing the three color channels and an image resolution of 224x224 pixels. ResNet50 systematically extracts features from the image through four convolutional blocks, each consisting of multiple layers, progressing from simple to more complex features. The extracted features are then decreased using a pooling layer that stores 2048 features. Finally, the last layer employs a softmax activation function, producing the probability that a sample matches one of the defined classes.

Deep learning based iris recognition systems can be designed in two distinct scenarios: closed-world and open-world. In a closed-world setup, the recognition system assumes a well-defined and known set of individuals for both training and testing. Since it operates under the premise that all encountered subjects are part of the existing dataset, it classifies them into one of the existing classes. Therefore, closed-world scenario suitable for controlled environments with a fixed and identifiable user base. On the contrary, an open-world setup is designed to address real-world scenarios where the system may encounter individuals not present in the training set. This setup is characterized by a broader applicability, enabling the recognition system to handle unknown subjects or impostors not accounted for during training. Therefore, we also generated results without using classification and softmax layers by exposing the 2048 features to calculate metrics between different subjects such as Euclidean distance.

Transfer learning is a method by which to speed up the convergence of a model by training a preexisting model with weights that are optimized for a similar task. To facilitate transfer learning,

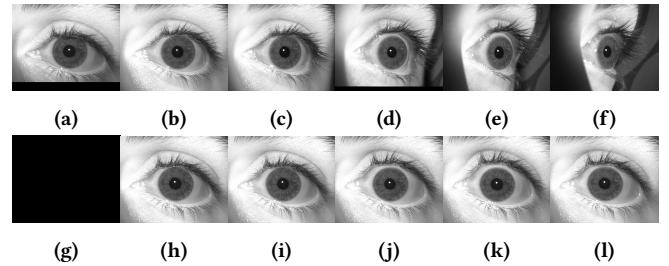


Figure 2: Example Orbital Camera and Frontal Camera Eye Images

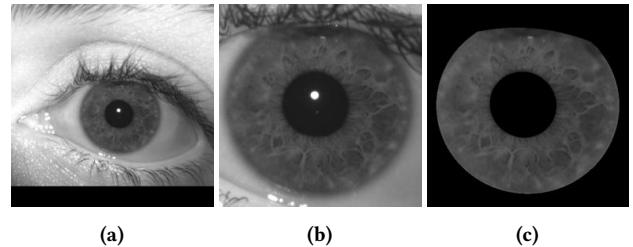


Figure 3: Example Periocular, Ocular, and Ocular Iris Images

ResNet50 was employed with pre-trained weights obtained from training on the ImageNet dataset [2]. The original output layer with 1000 classes that correspond with various real-world objects was replaced with a new softmax layer featuring 101 classes, aligning with the number of classes in our dataset. The flowchart in Figure 1 visually illustrates this process. Our implementation utilized the Python [11] library, PyTorch [8], along with Torchvision for creating, training, and validating the ResNet50 model.

2.2 Dataset

There is no publicly available iris dataset that includes images at different blink levels to evaluate the effects of blinking on deep learning iris recognition. To generate an iris dataset with blinks, we modified our existing off-angle iris dataset. The original dataset includes the left and right eye iris images from 113 subjects. Each subject was captured with two cameras, an orbital camera and a frontal camera. The orbital camera would rotate around the subject as the frontal camera remained stationary. 10 images per angle were taken by the orbital camera at angles -50 through 50, with a step size of 10 degrees. This resulted in 110 images per subject from the orbital camera. The frontal camera took images in sync with the orbital camera, resulting in a frontal version of each image for each off-angled image taken. However, the orbital camera occludes the frontal camera when at 0 degrees, resulting in 100 images per subject for the frontal camera. Figure 2 shows example images taken at 0 through 50 degrees by both cameras for the same subject. The first row displays images taken by the orbital camera, while the second row showcases those captured by the frontal camera at the same time as the orbital camera.

Once the images were captured, further modifications were done using image segmentation. Three pertinent cohorts were created:

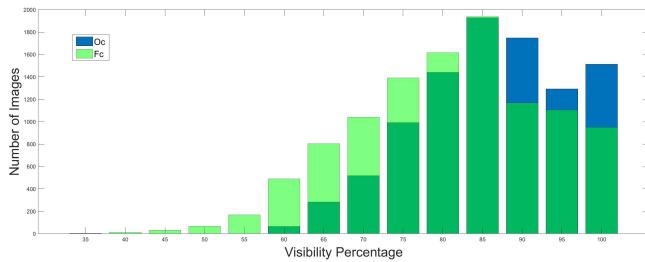


Figure 4: Original Ocular Iris Dataset Occlusion Distribution

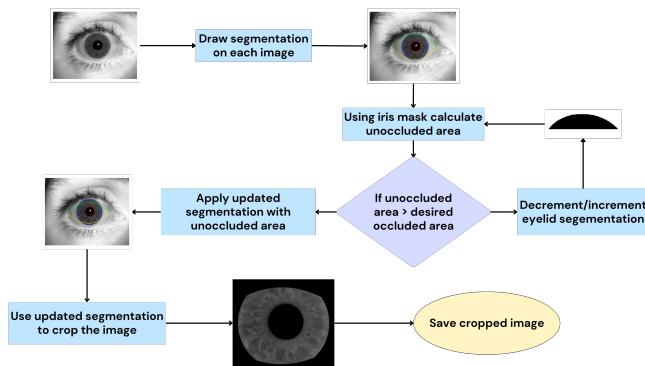


Figure 5: Flowchart of Synthetic Blink Occlusion Image

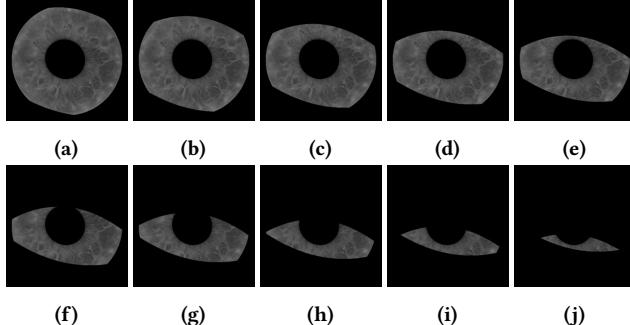


Figure 6: Examples of Synthetic Blink Occlusion Images

periocular, ocular, and ocular iris as shown in Figure 3. The periocular image on the left is cropped from the sides of the rectangular original image to generate a square image. It includes the most surrounding eye features in addition to the iris itself. The ocular cohort in the middle constrains the image to the iris and the features in its immediate vicinity. We generated this by further cropping the periocular image with the iris as the center point. Finally, the ocular iris cohort contains only the iris texture and was made by segmenting the iris's inner and outer boundaries and masking all other structures.

To design experiments to investigate the impact of blinking on iris recognition, the dataset needed to contain discrete cohorts of iris images at different levels of blinking. However, our original

dataset was not equipped for this purpose. Figure 4 shows the distribution of iris visibility in the original dataset changing by eye blinking. The light green plot shows the visibility distribution for the frontal camera and the blue plot shows the distribution of the orbital camera. The dark green part shows the overlap between the two cameras. In the original dataset, the majority of images had irises that were less than 20% occluded. Since this distribution was not conducive to measuring the effects of occlusion on iris recognition, the existing dataset had to be further modified by synthetically occluding existing iris images to create discrete blinking cohorts. This allowed for the use of our established database that has been used for previous experiments. Furthermore, synthetic occlusion does a good job of matching real-world blinking but allows for the precise measurement of its impact. Extending the ocular iris cohort was the obvious choice for this since ocular iris images have an existing eyelid mask that can be adjusted.

Figure 5 shows the flowchart describing the process of generating synthetically occluded ocular iris images. To create the synthetically occluded ocular iris images, the iris was segmented using elliptical inner and outer boundaries. Then, the eyelid was segmented using second-degree polynomial curves for the upper and lower eyelids. After iris and eyelid segmentation, the visibility area was calculated. This was achieved by taking the difference between the area of the iris ellipse and the visible area of the iris after masking was applied. Then, the visibility area was used to determine how occluded the iris was already, and how far it needed to be further occluded to fit into a designated cohort. Finally, the original iris image was segmented using the iris boundaries and the adjusted eyelid parameters, resulting in the synthetically occluded image.

The generated ocular iris images were sorted into 10 occlusion cohorts. Each cohort included 0-2%, 10%, 20%, ..., and 90% occlusion respectively. We chose to have ten cohorts and a step size of 10 percent as it is a good balance between having occlusion levels that are distinct, but not too far apart. To fill the cohorts, any images from the original dataset that did not specifically align with an occlusion cohort first had their mask adjusted to fit the occlusion level. For example, an iris image measured to be 25% occluded would be artificially occluded by another 5% to reach 30% occlusion. The 0% occlusion cohort included images ranging from 0-2% occlusion because there were very few iris images that contained irises with true 0% occlusion, and so 0-2% was used instead. Next, the eyelid mask was adjusted repeatedly to fill each of the corresponding occlusion cohorts. Figure 6 shows the results of sorting an ocular iris image into the 0-2% occlusion cohort and then incrementing the occlusion by 10% for each subsequent cohort.

Two separate datasets were created, one for the frontal camera images and one for the orbital camera images. Each dataset contained 101 subjects as some subjects had images unfit for the occlusion dataset. Furthermore, each dataset only contained images of the left eye. Finally, the generated dataset is not perfect. There are fewer images for the low occlusion cohorts, particularly the 0-2% cohort. This is because the original distribution of iris occlusion was not comprised of only iris images with 0% occlusion. Therefore, some images were not available for low occlusion cohorts. Table 1 shows the image count for each occlusion cohort generated for the frontal and orbital camera datasets. The original ocular iris dataset

Table 1: Image Counts of Each Occlusion Cohort

Cohort	Frontal Camera	Orbital Camera
Original	9,880	10,886
0-2% Occl.	1,382	831
10% Occl.	3,614	911
20% Occl.	7,360	6,024
30% Occl.	9,310	8,799
40% Occl.	9,868	10,415
50% Occl.	9,876	10,811
60% Occl.	9,878	10,882
70% Occl.	9,880	10,886
80% Occl.	9,880	10,001
90% Occl.	9,880	10,886

image count is also shown for the frontal camera dataset because it was used during the experimentation.

For usage with ResNet50, we processed the images before feeding them into the network. The images for the training and validation datasets were normalized to the ImageNet standard [2]; the mean values being 0.485, 0.456, and 0.406, and the standard deviation values being 0.229, 0.224, and 0.225 for the red, green, and blue channels respectively. The images were also resized to a 224x224 resolution image for the ResNet50 input. For the training dataset, some augmentations were applied before the normalization and resizing. These included a random -10% to 10% translation of the image, a random -20 to 20-degree image rotation using bicubic interpolation, and a Gaussian blur with a kernel size of 3 and sigma of 0.1 minimum to 0.5 maximum. These augmentations were shown to improve results and reduce overfitting during the training process.

2.3 Experimental Design

Our goals in the experiments were twofold: i) to measure the effects of blinking on a baseline model without much exposure to blinked images and ii) to investigate potential performance improvements using synthetically occluded images during training. For these reasons, different models were trained and validated using the same basic experimental setup.

The training dataset for each model used frontal camera images taken at all angles except angle 0, meaning -50, ..., -10, 10, ..., and 50 degrees. The testing and validation dataset used only angle 0 images from the orbital camera. This means that training and validation were done using only images captured from the front. However, the validation images were taken at a different time than the training images and therefore had enough variance to avoid overfitting. The occlusion cohorts used in the training dataset were configured differently in each experiment. Meanwhile, every occlusion cohort, except for the original ocular iris dataset, was included in the testing and validation dataset. This was done to get an accurate testing loss during training and to evaluate each occlusion cohort distinctly during validation.

ResNet50 training was performed with the following parameters: The cross-entropy loss was used as a loss function and mini-batch gradient descent was the optimizer with the following parameters: a batch size of 128, a learning rate of 0.01, a momentum of 0.9, and

Table 2: Model Training Metadata

Model	Image Count
Original Dataset	9880
Original + 0-30% Occl.	31546
Original + 0-60% Occl.	61168
Original + 0-90% Occl.	90808

Table 3: Validation Dataset Metadata

Cohort	Image Count
0-2% Occl.	80
10% Occl.	104
20% Occl.	561
30% Occl.	801
40% Occl.	979
50% Occl.	988
60% Occl.	990
70% Occl.	990
80% Occl.	904
90% Occl.	990

a weight decay rate of 2e-5. For each epoch in training, each batch in the training dataset was passed through the network and the weights were adjusted using the batch loss. The batch losses were summed and divided by the number of batches to get the mean training loss for the epoch. Next, the testing dataset was passed through the network in batches, and the mean testing loss was computed in the same manner as the mean training loss. Training was halted when the testing loss stopped improving.

Using the aforementioned methodologies, four experiments were designed using different training datasets: 1) only the original ocular iris dataset for training, 2) the original ocular iris dataset and the 0-2% through 30% occlusion cohorts, 3) the original ocular iris dataset and the 0-2% through 60% occlusion cohorts, and 4) the original ocular iris dataset and the 0-2% through 90% occlusion cohorts. The first experiment establishes a baseline performance for a model trained on mostly open ocular iris images, while the subsequent experiments investigate the performance changes caused by training with increasingly higher occlusion.

3 RESULTS

This section presents the performance results of four sets of experiments mentioned in the experimental design. To evaluate the models in each experiment, we used two different metrics including the receiver operating characteristic (ROC) curve and Euclidean distance histogram plots. We chose these two metrics due to their extensive use in previous iris biometric research. In addition, these metrics provide important information about the model's performance in both closed-world and open-world scenarios. Different ROC plots were generated for each occlusion cohort in the validation dataset, resulting in a stratified view of the model performance for each level of occlusion. Euclidean distance histograms were generated for the 0%, 30%, 60%, and 90% occlusion cohorts to show

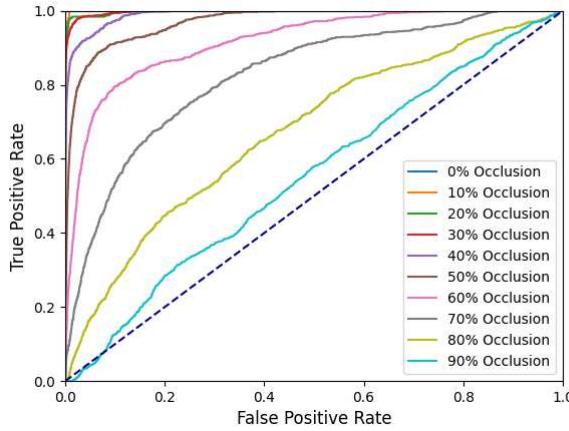


Figure 7: ROC Plots for Baseline Results

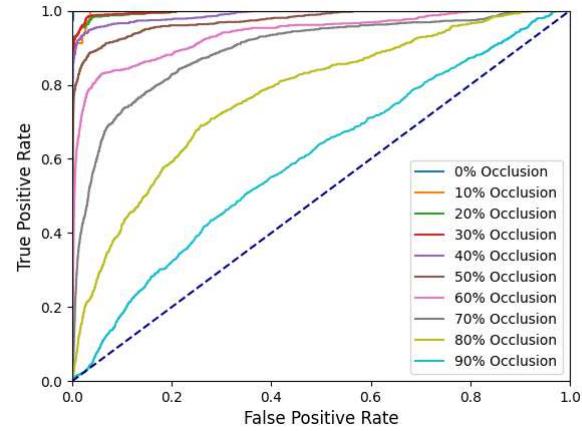


Figure 9: ROC Plots for Original and 0-30% Occlusion Cohorts

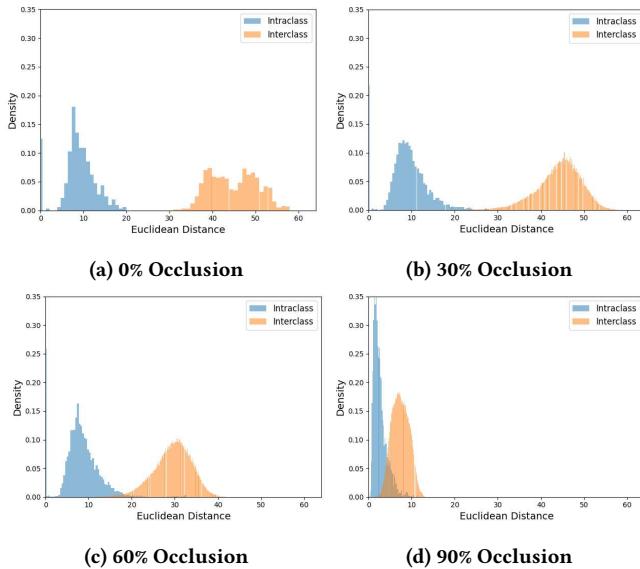


Figure 8: Euclidean Distance Histograms of Baseline Model for 0-2%, 30%, 60%, and %90 Occlusion Cohorts

how the feature distances between irises changed at different levels of occlusion. Furthermore, the intra-class and inter-class images were plotted on separate curves to show the model's ability to distinguish between the features of a single subject and the rest, an important metric for open-world applications. Table 2 shows the number of images used during the training of each model and Table 3 shows the number of images used in each occlusion cohort to validate each model. We can see that the number of available images at lower occlusions was small compared to higher occlusions. Therefore, some subjects were not equally represented at low levels of occlusion.

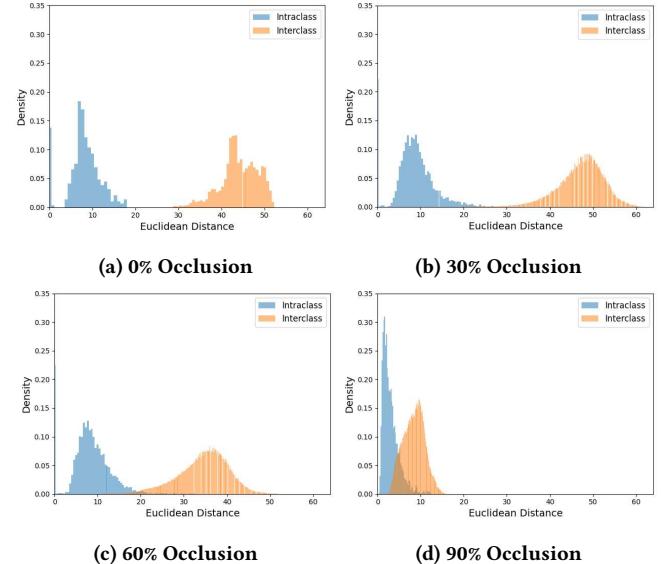


Figure 10: Euclidean Distance Histograms of Original and 0-30% Occlusion Model for 0-2%, 30%, 60%, and %90 Occlusions

3.1 Baseline Results

In the first set of experiments, we trained a model using only the original ocular iris dataset to determine the baseline performance. ROC curves shown in Figure 7 presents the performance at 0-2%, 10%, ..., and 90% iris occlusion cohorts where the x-axis is the false positive rate and the y-axis is the true positive rate. It shows a steady reduction in performance as the occlusion increases. We can also see that the performance does not drop severely until the occlusion surpasses 40%. This makes sense as the distribution of occlusion in the original dataset includes many images with 0-30% occlusion.

Figure 8 shows the Euclidean Distance histograms of the baseline model for different occlusion cohorts. It presents how the distances

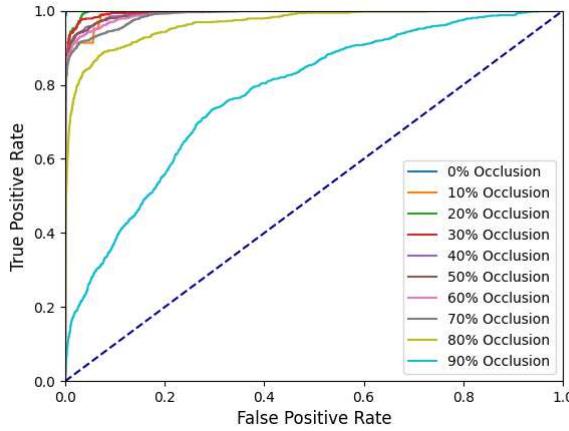


Figure 11: ROC Plots for Original and 0-60% Occlusions

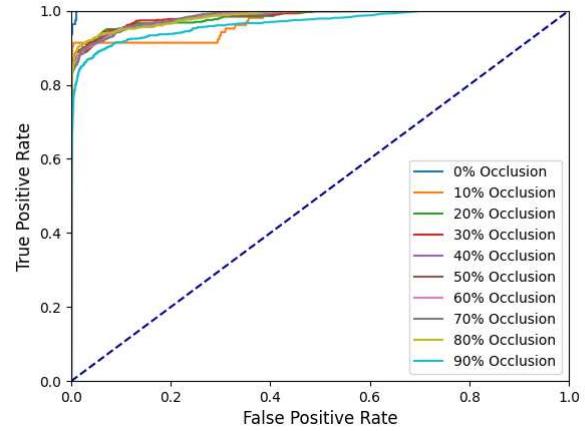


Figure 13: ROC Plots for Original and 0-90% Occlusions

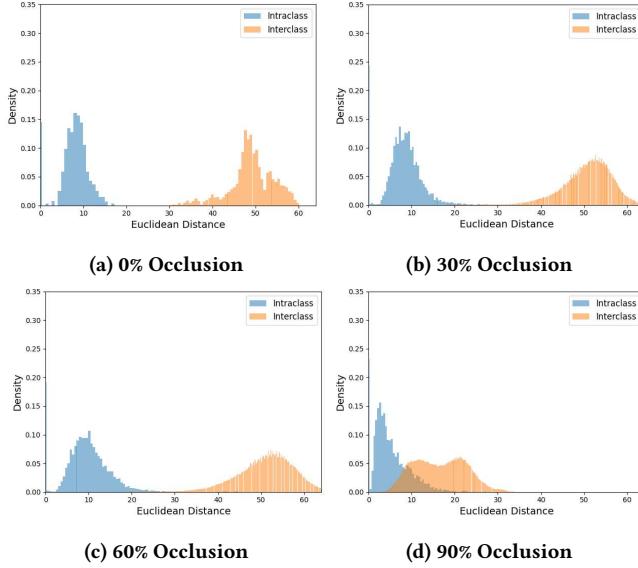


Figure 12: Euclidean Distance Histograms of Original and 0-60% Occlusion Model for 0-2%, 30%, 60%, and %90 Occlusions

between the extracted features change based on the occlusion. We can see that at low occlusion, the intra-class images and inter-class images have very different features. However, as occlusion increases, the variance between the features decreases, eventually leading to severe overlap between the intra-class and inter-class distances. This is because the masking of the iris obscures features that the model is adept at extracting, resulting in much of the extracted feature vector being flat.

3.2 Results for Original and 0-30% Occlusions

In the second set of experiments, we trained the model with the original and 0-30% occlusion cohorts. It showed slightly improved results compared to the baseline experiment. Figure 9 shows that the

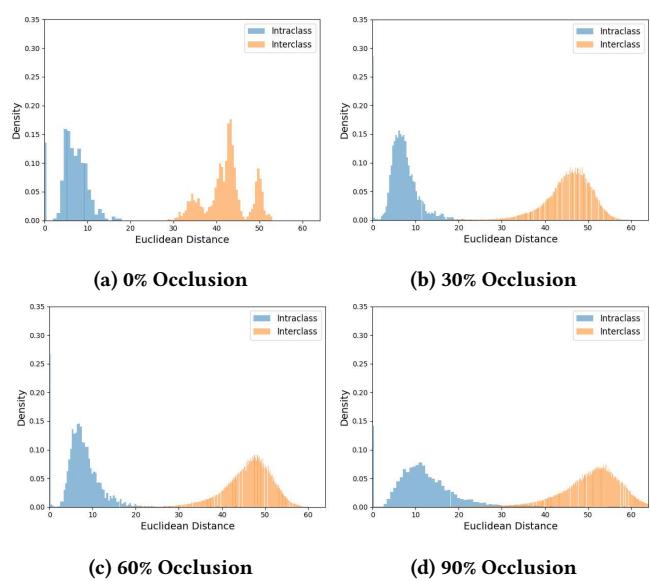


Figure 14: Euclidean Distance Histograms of Original and 0-90% Occlusion Model for 0-2%, 30%, 60%, and %90 Occlusions

performance using ROC plots where it starts to decrease severely after 50% occlusion. The difference between this results and the baseline results is small because the images in the original ocular iris dataset overlaps substantially with the 0-2% through 30% occlusion cohorts.

Figure 10 shows the distribution histograms of the Euclidean Distance from the second model for different occlusion cohorts due to the eye blinks. We observed that the distribution of Euclidean distances is mostly unchanged when compared to the first model, which makes sense considering the overlap between the images in the original ocular iris dataset and the 0-2% through 30% occlusion cohort images.

3.3 Results for Original and 0-60% Occlusions

The third experiment trained the model with the original and 0-60% occlusion iris image cohorts, showing significant performance improvements when compared to the baseline and second set of experiments. Figure 11 shows the performance results using ROC plots. We observed that there is little performance degradation until the occlusion level surpasses 70%.

The Euclidean distance histograms for the third set of experiments are shown in Figure 12. It shows some distinct changes where we observed an overall greater separation of intra-class and inter-class distances. This includes the 90% occlusion cohort, where we can see more separation between the distances compared with the first two sets of experiments.

3.4 Results for Original and 0-90% Occlusions

In the final set of experiments, we trained the model with the original and 0-90% occlusion iris image cohorts. The fourth model shows a consistent, but slightly inferior performance for any level of occlusion. We observed that each of the plots, excluding the 0% occlusion cohort follow a similar curve. However, compared to the previous models, the curve is not as tightly bound to the upper left of the ROC, indicating a sacrifice in performance.

The fourth set of Euclidean distance histograms shows a distinct separation of distances between intra-class and inter-class images for all occlusion levels. In Figure 14, we can see that, unlike for previous models, the 90% occlusion cohort's intra-class and inter-class distances have high separation.

4 DISCUSSION

Our results draw a conclusive trend, that model performance increases as the level of occlusion shown during training increases. However, there is some nuance. When training on all levels of occlusion, including 90%, the overall performance of the model was slightly hindered. This is due to a limitation of the ocular iris segmentation. All features of the image are excluded except for the iris, meaning that at high occlusion, there are very few unmasked pixels to be observed as distinct features. Due to this limitation, ResNet50 learned to prioritize features in the lower-middle iris. The result of this can be seen in the Euclidean distance histograms, particularly for the third model and fourth model, where there is still a high separation between intra-class and inter-class distances despite limited iris features. ResNet50 prioritizing only some of the total available features in the iris images caused the performance for low occlusion images to degrade as other viable features were ignored. For this reason, we can say that, for the ocular iris dataset, it is likely best to expose up to 60% or 70% occlusion to the model during training.

The success in separating the features between intra-class and inter-class features indicates that training the model to account for blinking is a viable way to augment biometric systems for open-world applications as the limited number of features still have a unique profile. Of course, this could change given thousands or millions of subjects, where there is bound to be more overlap for a small subsection of features. Recognizing our dataset limitations, future research might benefit from employing eye images where features surrounding the iris are not masked, mitigating ocular iris

image constraints and potentially enhancing performance through the inclusion of additional distinctive features beyond the iris.

It is also worth discussing the effects of fewer available images at lower occlusions. One distinctive result is the 0-2% occlusion validation cohort displaying almost perfect performance for every model. This may be caused by the limited number of validation images at that level of occlusion; in this case, only 80. Those limited images may be optimal for each subject, with high feature variance between subjects, resulting in an almost perfect performance. Despite this outlier, the results are still quite clear.

5 CONCLUSION

This paper aimed to extend the understanding of CNN-based stand-off iris biometric systems by examining the impact of blinking on iris recognition performance. The methodology involved simulating blinking on ocular iris images through masking and subsequently retraining the ResNet50 model via transfer learning. Initially, a baseline model was established using the original ocular iris dataset, followed by experimentation with augmented training employing progressively more occluded images. The baseline model exhibited a linear decrease in performance as occlusion levels increased. In contrast, models trained with synthetically occluded images displayed substantial performance enhancements. However, challenges emerged when training with very high occlusion, highlighting limitations associated with ocular iris images due to the constrained availability of features for the model. A noteworthy recommendation for future research is to explore augmentation techniques involving unmasked blinked images and delve into alternative forms of augmentation tailored for iris recognition training.

ACKNOWLEDGMENTS

This project was made possible by support from Secure and Trustworthy Cyberspace (SaTC) program of The National Science Foundation (NSF) under grant award CNS-1909276 and CNS-2100483.

REFERENCES

- [1] John Daugman. 2007. New Methods in Iris Recognition. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)* 37, 5 (2007), 1167–1175. <https://doi.org/10.1109/TSMCB.2007.903540>
- [2] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. Imagenet: A large-scale Hierarchical Image Database. In *2009 IEEE conference on computer vision and pattern recognition*. Ieee, 248–255.
- [3] Abdulrahman Aminu Ghali, Sapiee Jamel, Kamaruddin Mohamad, Nasir Yakub Abubakar, and Mustafa Mat Deris. 2017. A Review of Iris Recognition Algorithms. *JOIV: International Journal on Informatics Visualization* 1 (11 2017), 175. <https://doi.org/10.30630/joiv.1.4-2.62>
- [4] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep Residual Learning for Image Recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 770–778.
- [5] Mahmut Karakaya. 2021. Iris-ocular-periocular: Toward More Accurate Biometrics for Off-Angle Images. *Journal of Electronic Imaging* 30, 3 (2021), 033035. <https://doi.org/10.1117/1.JEI.30.3.033035>
- [6] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. ImageNet Classification with Deep Convolutional Neural Networks. In *Advances in Neural Information Processing Systems*, F. Pereira, C.J. Burges, L. Bottou, and K.Q. Weinberger (Eds.), Vol. 25. Curran Associates, Inc.
- [7] Peizhong Liu, Guo Jing, Szu-Han Tseng, Koksheik Wong, Jiann-Der Lee, Chen-Chieh Yao, and Daxin Zhu. 2017. Ocular Recognition for Blinking Eyes. *IEEE Transactions on Image Processing* PP (06 2017), 1–1.
- [8] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith

Chintala. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems 32*. Curran Associates, Inc., 8024–8035.

[9] Stephanie A. C. Schuckers, Natalia A. Schmid, Aditya Abhyankar, Vivekanand Dorairaj, Christopher K. Boyce, and Lawrence A. Hornak. 2007. On Techniques for Angle Compensation in Nonideal Iris Recognition. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)* 37, 5 (2007), 1176–1190.

[10] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. 2015. Going Deeper with Convolutions. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 1–9.

[11] Guido Van Rossum and Fred L. Drake. 2009. *Python 3 Reference Manual*. CreateSpace, Scotts Valley, CA.