ImmerIris: A Large-Scale Dataset and Benchmark for Off-Axis and Unconstrained Iris Recognition in Immersive Applications

Supplementary Material

This supplementary material provides more technical and experimental details that were streamlined in the main text due to space limitations, and we hope that it will be of interest to our readers. It mainly includes:

(A) Summary of Existing Datasets;
(B) Annotation Details;
(C) Protocol Details;
(D) Additional Visualizations;
(E) Discussion of Normalization;
(F) Additional Experimental Results;
(G) Ethical Discussion;
(H) Dataset Availability.

Tables 1-7 and Figures 1-6 refer to items in the main text.

A. Summary of Existing Datasets

Table 2 in the main text compares a shortlist of iris recognition datasets that are most closely related to ImmerIris. Here, we provide an extended summary of them.

- Figure 7 illustrates sample images from existing datasets and ImmerIris. It is worth noting that existing datasets mostly acquire images on-axis, which are inadequate for capturing the perspective distortion caused by immersive acquisition. PolyU Iris DB [20] is similar to ImmerIris in that it also uses a VR device for acquisition. However, its camera is placed below the eyes, and thus the distortion is less severe. In contrast, ImmerIris places the camera at the corner, hence perspective distortion is pronounced.
- Table 8 provides an extended list of existing datasets. Note that ImmerIris is the largest in total image volume to date and among the largest in terms of subjects involved.

For further discussion of existing datasets, we refer readers to this survey [14].

B. Annotation Details

We provide further details on how the data annotation process (Sec. 3.2) is carried out.

We first sample and manually annotate 10K images. Specifically, for each image, a human expert determines a score between 0 and 1 regarding the severity of occlusion and the extremity of gaze, reflection and blur. For the extremity of dilation, the iris region is first segmented using an off-the-shelf pretrained model, and then an iris-to-ocular area ratio is computed, where a smaller ratio represents a more extensively dilated pupil. We use these samples and annotations to train a VGG-based model, which is then used to automatically annotate the remaining images. The thresholds are set empirically by the human expert. Images with a quality score below the threshold are denoted as degraded in their corresponding quality dimension. Table 9 reports the thresholds and the number of samples degraded exclusively in each single dimension.

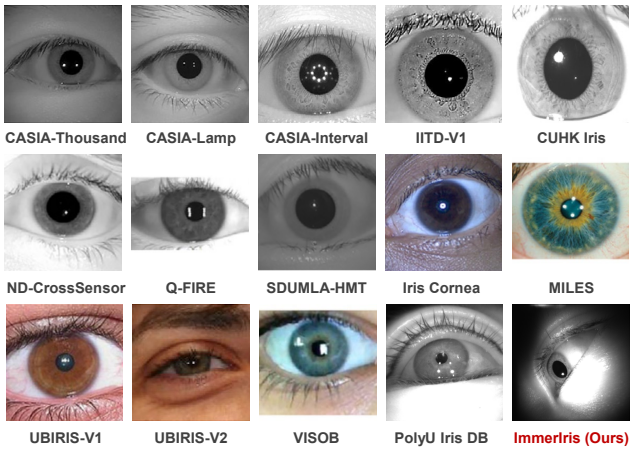

Figure 7. Comparison of samples images from existing datasets and ImmerIris. ImmerIris best captures perspective distortion.

Table 8. Comparison of existing iris recognition datasets and ImmerIris in terms of acquisition setup and scale.

Dataset	Axis	Conditions	# Img (Subject)
CASIA-IrisV1 [4]	On	Controlled	36,240 (50)
CASIA-IrisV4 [5]	On	Controlled	20,000 (1000)
IITD-V1 [12]	On	Controlled	1,120 (112)
CUHK Iris [6]	On	Controlled	254 (18)
ND-CrossSensor [23]	On	Controlled	117,503 (676)
WVU Twins Day [15]	On	Controlled	N/A (76)
ND-TimeLapse [3]	On	Controlled	6797 (23)
Q-FIRE [10]	On	Controlled	1800 (195)
SDUMLA-HMT [24]	On	Controlled	1600 (106)
Iris Cornea [11]	On	Controlled	780 (39)
MILES [9]	On	Controlled	832 (25)
UPOL [8]	On	Controlled	354 (64)
UBIRIS.V1 [16]	On	Semi-Ctrl	1,877 (241)
UBIRIS.V2 [17]	On	Semi-Ctrl	11,102 (261)
VISOB [18]	On	Semi-Ctrl	95,107 (550)
CASIA-BTAS [25]	On	Semi-Ctrl	4,500 (150)
CASIA-Iris-M1 [26]	On	Semi-Ctrl	11,000 (630)
PolyU Iris DB [20]	Off	Semi-Ctrl	142,005 (384)
ImmerIris (ours)	Off	Unconstrained	499,791 (564)

C. Protocol Details

We describe the details of how we sample images according to the protocol designs and formulate the evaluation.

Image selection. For each protocol, we first select all images that meet the protocol definition to form a "pool" (*e.g.*, 6660 left-eye images with occlusion and without other degradations form the pool for Immer-Occlusion; 40,506 normal-quality right-eye images form the pool for Immer-Control, on their respective eyes). In Sec. 3.3, we have selected 166 subjects for the overall test set. Occasionally, it occurs that a subject has no image included in a

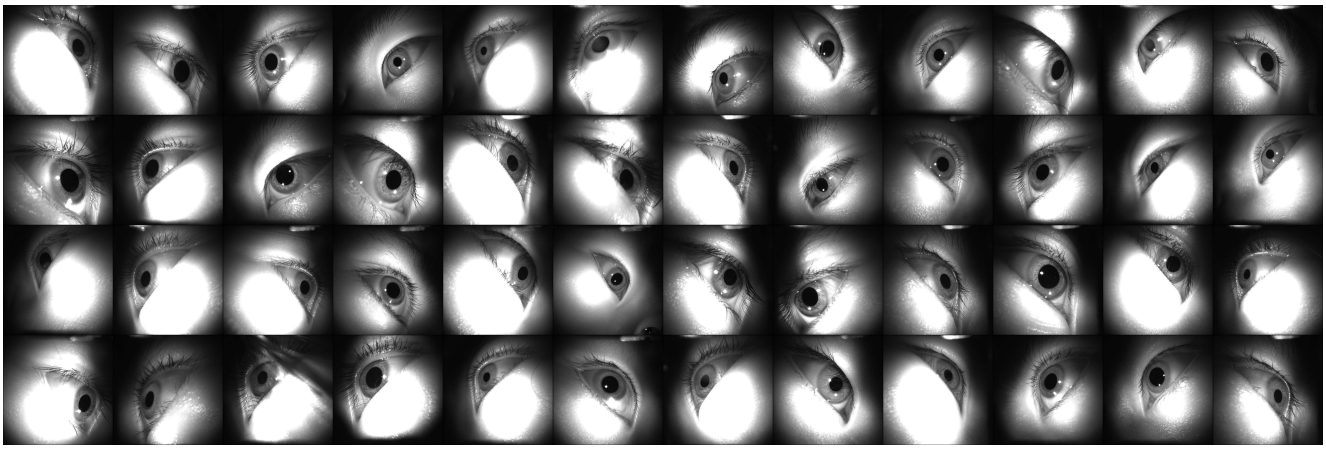

Figure 8. Normal-quality samples of the proposed ImmerIris dataset, in addition to Fig. 4(b).

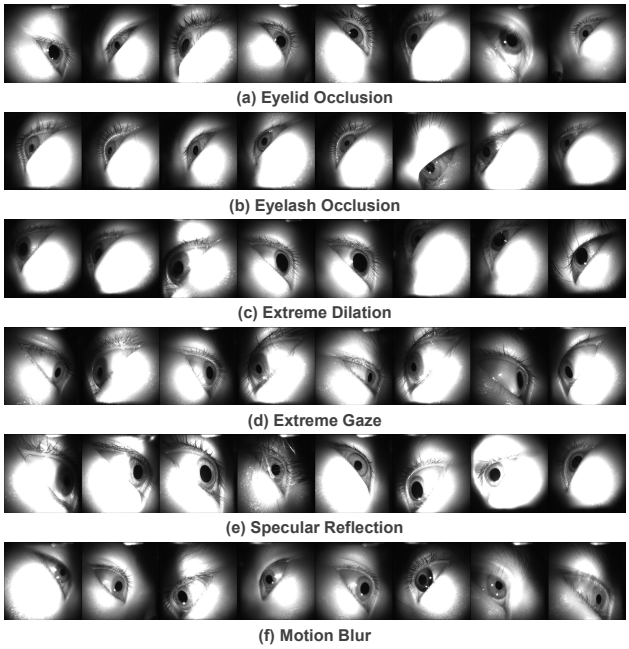

Figure 9. Degraded samples in addition to Fig. 5(b).

Table 9. Thresholds and number of degraded samples (training and test sets combined) of each quality dimension.

Quality Dimension	Threshold	#Samples
Eyelid Occlusion	0.6	46716
Eyelash Occlusion	0.3	3664
Extreme Dilation	0.6	10703
Extreme Gaze	0.59	31007
Specular Reflection	0.93	3361
Motion Blur	0.36	17105

pool (*e.g.*, many subjects may have no extremely dilated images). In such cases, the pool (and hence the corresponding protocol) can cover fewer than 166 subjects.

Pair formation. We then uniformly sample pairs from the pool according to each protocol's design. Some protocols share the same pool (*e.g.*, Gaze and Control both use all normal-quality samples), but their pair formation differs significantly due to the protocol specifications. In principle, we traverse all feasible pairs, but the total number may be excessive for certain protocols (*e.g.*, up to 19.8M impostor pairs for Gaze), which would drastically slow down evaluation. Therefore, we set a maximum capacity for each protocol. For verification, we sample at most 1.5M genuine pairs per protocol, and 2M or 3M impostor pairs for protocols studying isolated factors or combined factors, respectively. For identification, we choose at most 100 probe images per class. Samples are randomly shuffled before pairing, and we thus expect the sampled pairs to be uniformly distributed across the dataset rather than limited to certain subjects.

Pair sampling and protocol scale. For verification, we report FRR@FAR(α). By definition, FRR@FAR requires at least $\frac{1}{\alpha}$ impostor pairs to achieve sufficient precision at the FAR=α threshold (*e.g.*, at least 100K pairs for α=1e-5). All protocols except Dilation have abundant (>2M) impostor pairs. The Dilation protocol has fewer due to its specification (*i.e.*, requiring both images to be dilated), yet it still has at least 121K pairs, being sufficient for stable estimation at that precision. The number of genuine pairs also varies across protocols. In particular, Occlusion, Dilation, and Light have relatively fewer genuine pairs. In theory, this may slightly affect robustness; however, since even the smallest counts exceed the minimum required scale by a significant margin, the practical impact is very limited.

Illumination as a control variable. We do not explicitly include illumination as an independent control variable in our protocols (Tab. 2). Instead, we employ the Light protocol to highlight the divergence between the most and least illuminated samples. This is because controlling both 11 illumination levels and 9 gaze directions would drastically reduce the number of available pairs, thereby limiting the maximum achievable precision under the protocol.

Dual-eye specification. For the dual-eye protocols, we require that each bi-ocular sample (*i.e.*, a pair of left and right eye images, considered as one sample) be captured from the same subject at the same moment. Therefore, no cross-subject or cross-time bi-ocular combinations are included. In protocols involving degradation (*e.g.*, Occlusion), a bi-ocular sample is enrolled in the pool

Table 10. Protocol scales for single-eye and dual-eye verification. "#ID, #Sample, #G-Pair, #I-Pair" denote the number of subjects, samples, genuine pairs, and impostor pairs, respectively.

Protocol	Left Eye				Right Eye				Dual Eyes			
	#ID	#Sample	#G-Pair	#I-Pair	#ID	#Sample	#G-Pair	#I-Pair	#ID	#Sample	#G-Pair	#I-Pair
Occlusion	146	6660	75.8K	2M	138	9297	124K	2M	154	13020	181K	2M
Dilation	46	1856	22.6K	177K	52	1396	14.8K	121K	66	2557	29.5K	357K
Light	163	11160	19K	2M	158	9690	16.6K	2M	157	8140	14K	2M
Gaze	160	47376	1.5M	2M	156	40506	1.5M	2M	138	31029	1.5M	2M
Control	160	47376	1.04M	3M	156	40506	880K	3M	138	31209	1.24M	3M
Fix	166	77040	1.5M	3M	164	74824	1.5M	3M	164	72090	1.5M	3M
Select	166	51691	1.5M	3M	164	50356	1.5M	3M	-	-	-	-
Any	166	77040	1.5M	3M	163	74824	1.5M	3M	164	72090	1.5M	3M

as long as either eye exhibits the respective challenging factor. In such cases, the number of bi-ocular samples is substantially larger than that in either single-eye protocol, yet still smaller than their combined total. In contrast, for protocols excluding degradation (*e.g.*, Gaze), both eyes are required to be normal-quality samples. This results in fewer bi-ocular samples compared to the single-eye settings.

Table 10 reports the protocol scales for single-eye and dual-eye verification.

D. Additional Visualizations

We provide more sample images to illustrate the proposed ImmerIris dataset. All images are selected at random.

- Figure 8 shows additional samples of normal quality, as an extension of Fig. 4(b) of the main text.
- Figure 9 presents more samples from each quality degradation category, namely eyelid occlusion, eyelash occlusion, extreme dilation, extreme gaze direction, specular reflection, and motion blur, as an extension of Fig. 5(b). Note that one sample may exhibit multiple degradations.
- Figure 10 shows ocular images from both eyes acquired at different gaze points (*i.e.*, red squares numbered 1-9) in a continuous session from the same person. Intra-subject variation in gaze direction can be observed.

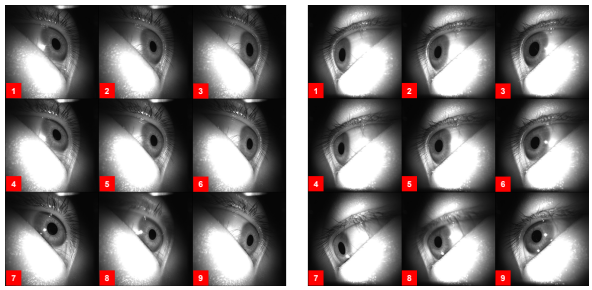

Figure 10. Ocular images from both eyes at different gaze points acquired from the same subject. Intra-subject variation occurs.

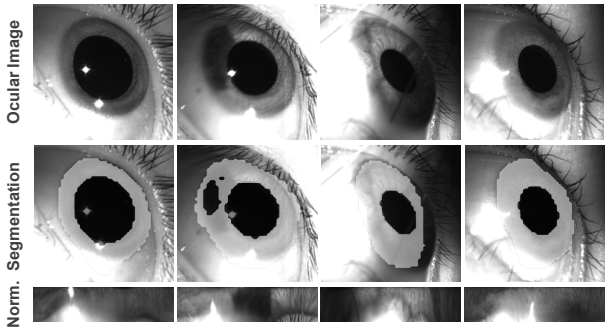

Figure 11. Normalization failure cases for off-axis images.

E. Discussion of Normalization

We here discuss why normalization can yield unsatisfactory iris textures in immersive iris recognition.

Normalization [7] (also referred to as rubber-sheet or polar unwrapping in some literature) in traditional iris recognition implicitly assumes a near-on-axis imaging geometry. It can be subdivided into three phases, namely: *Segmentation*, where the iris region is detected and extracted from the ocular image; *Contour parameterization*, where circles or ellipses are fitted to the iris boundaries to obtain a geometric representation; *Polar transformation*, where the fitted circles or ellipses are unwrapped into a rectangular strip of normalized texture. In immersive iris recognition, where images are acquired off-axis, several assumptions underlying each phase are violated, leading to accumulated errors.

Error in segmentation. During segmentation, the off-axis view increases reflections and occlusions, making boundary detection unstable. Under perspective projection, the iris region is no longer a circular disk; instead, it becomes a projected ellipse whose axes do not correspond to the true iris radii, and whose pupil and iris ellipses are no longer coaxial. Segmentation therefore introduces a systematic geometric error that is inherited in every later phase.

Error in contour parameterization. Even with perfect segmentation, contour parameterization can introduce a mismatch. Rubber-sheet normalization assumes the pupil and iris boundaries can be parameterized as either two circles or two *coaxial* ellipses. However, under off-axis perspective, the pupil and iris do not share the same center, and their shape distortion depends on the camera pose and the true 3D eye geometry. Therefore, even a perfect ellipse fit cannot recover the true geodesic structure of the iris surface, and normalization based on circular or coaxial-elliptical geometry inevitably produces non-uniform stretching and compression in the normalized texture.

Error in polar normalization. Polar normalization typically uses

Table 11. Identification rank-1 accuracy (↑) of SOTAs and the proposed method on different evaluation protocols.

Eye	Method	Control	Fix	Select	Any	Occlusion	Light	Gaze
Left	Gabor [7]	91.99	86.44	48.52	49.40	65.72	65.57	50.01
	OM [19]	85.55	79.04	45.15	44.79	45.58	62.30	47.16
	Maxout [26]	97.70	93.97	76.97	76.12	77.39	81.97	79.96
	Maxout-BA [21]	98.74	95.79	82.18	81.78	89.40	88.52	85.75
	UE-UGCL [21]	99.35	97.27	89.92	88.83	93.64	96.72	93.25
	CM [22]	99.59	98.19	93.59	92.65	95.41	95.08	**96.14**
	ComplexIrisNet [13]	99.67	97.89	92.52	91.14	93.64	96.72	95.18
	NormKeep	**99.76**	98.26	92.28	91.99	95.41	98.36	94.75
	NormFree (ours)	99.52	**98.91**	**93.87**	**94.39**	**98.23**	**98.36**	95.49
Right	Gabor [7]	86.80	80.41	45.10	45.63	71.43	59.21	46.64
	OM [19]	84.01	77.20	44.06	44.77	62.86	55.26	46.26
	Maxout [26]	93.72	89.91	70.41	70.34	86.26	73.68	72.85
	Maxout-BA [21]	95.78	92.21	76.49	75.65	88.71	72.37	79.08
	UE-UGCL [21]	98.21	95.60	86.58	86.41	95.10	76.32	89.28
	CM [22]	99.18	97.49	91.90	91.47	96.60	97.37	94.28
	ComplexIrisNet [13]	99.30	97.07	91.00	90.04	96.46	92.11	94.00
	NormKeep	99.25	97.76	91.55	91.43	98.50	90.79	93.86
	NormFree (ours)	**99.51**	**98.86**	**94.44**	**94.17**	**99.57**	**100.00**	**94.44**
Dual	Gabor [7]	76.56	73.23	-	30.57	58.06	43.08	40.75
	OM [19]	64.68	61.34	-	22.83	46.93	16.92	32.42
	Maxout [26]	90.92	88.55	-	59.10	78.60	66.15	70.32
	Maxout-BA [21]	93.48	91.37	-	65.86	83.74	72.31	76.20
	UE-UGCL [21]	98.01	96.25	-	79.34	95.72	73.85	89.04
	CM [22]	97.97	96.54	-	86.03	94.29	80.00	93.54
	ComplexIrisNet [13]	98.49	96.98	-	84.72	94.58	86.15	92.94
	NormKeep	98.33	96.98	-	85.26	96.86	83.08	92.30
	NormFree (ours)	**99.03**	**98.27**	-	**88.93**	**97.48**	**98.41**	**93.84**

$$(r, \theta) \rightarrow (x(r, \theta), y(r, \theta)) \tag{1}$$

where θ is uniformly sampled and r is interpolated between the pupil and iris boundaries. This assumes that 1) constant angular spacing corresponds to equal arc length, 2) radial lines correspond to true anatomical meridians, and 3) the distance from the pupil to the iris boundary corresponds to the true radial dimension of the texture. None of these assumptions hold precisely in off-axis conditions. Hence, polar unwrapping imposes an incorrect parameterization, leading to template drift (*i.e.*, different iris textures for the same person) and thus poor matching performance.

Though some SOTAs attempt to improve normalization design or incorporate post-processing, they remain largely non-intuitive and suboptimal. For example, an adaptive normalization strategy [20] is found to improve recognition performance only marginally compared with our normalization-free paradigm, as previously shown in Tab. 7. Hence, we believe that being normalization-free reveals a promising direction for methodological advance.

Visualization. Figure 11 illustrates several failure cases of normalization when applied to off-axis ocular images. Each column (top to bottom) shows an original off-axis image, its segmentation result, and the corresponding normalized texture. Errors occur in segmentation, where reflections are mapped as bright vertical bands and eyelid or eyelash regions are incorrectly unwrapped as iris texture. Errors also occur in the normalized textures, which exhibit non-uniform stretching and compression, missing texture patterns, and strong geometric skew.

Table 12. Verification FRR@FAR (↓) of NormKeep trained and tested on on-axis CASIA-T (C) and off-axis Immer-Any (I), covering all 4 train-test combinations.

Method	Train-Test	Performance		
	FRR@FAR (%) ↓	1e-1	1e-3	1e-5
NormKeep	I / I	5.84	29.72	55.23
	C / C	0.50	2.59	6.41
	C / I	32.88	76.80	91.14
	I / C	2.64	14.51	39.83

F. Additional Experimental Results

Implementation Details. This is the implementation details of Sec. 5.1. We rerun SOTAs with available official codes [21, 22, 26] on our datasets and re-implement the remaining ones [7, 13, 19] to our best effort, while acknowledging potential inconsistencies. SOTAs are trained according to their recommended settings. Iris normalization is performed using OpenIris [1], unless otherwise specified by SOTAs.

Identification. We report the full performance in Tab. 11. It was streamlined in Tab. 6 due to space constraints. Overall, we observe similar trends as in the verification task: The protocols reveal distinct challenges from distortion, variation, and degradation in immersive iris recognition compared with the traditional setup, where SOTAs still show considerable room for improvement under these factors; Dispensing with normalization also achieves a considerable performance gain that is worth future investigation.

Cross-domain evaluation. We have learned in Tab. 3 that

normalization-based methods achieve success on-axis (when both trained and tested on CASIA-T) yet generalize to off-axis settings (trained on CASIA-T and tested on Immer-Any) unsatisfactorily. A natural question that arises is whether the performance gap originates mainly from Immer-Any being more challenging than CASIA-T, or simply from the domain gap between datasets.

To verify this, we train and test our normalization-free baseline (*i.e.*, NormKeep) on all four train-test dataset combinations. Table 12 reports their FRR@FAR, where the abbreviations indicate train-test datasets; *e.g.*, "I/C" indicates training on ImmerIris and testing on CASIA-T. Similar to previously observed in Tabs. 3 and 4, it performs well on-axis (C/C) and less satisfactorily off-axis (I-I). It further shows that a model trained on-axis cannot be generalized to off-axis (C/I), where performance drastically degrades. We highlight the opposite setting of C/I, where the model is trained off-axis and tested on-axis (I/C). It achieves lower performance than C/C due to the domain gap between two datasets, yet much *better* performance than I/I, which strongly indicates that the performance gap is attributed to the inherent challenges of the immersive scenario.

G. Ethical Discussion

This study involves the collection and analysis of iris images from human participants. Here, we discuss how the acquisition process meets ethical and privacy compliance requirements, as well as its potential societal impact.

Privacy compliance. We take responsible measures to ensure informed consent and data anonymization.

- **IRB approval.** All data acquisition procedures were conducted under the approval of the Institutional Review Board (IRB) of our institution.
- **Informed consent.** Prior to participation, each subject received a Q&A document clearly explaining the purpose of this research, the scope of data collection, the downstream processing protocols, and their privacy rights, including the right to withdraw at any time without penalty. To ensure fully informed participation, every subject signed two consent forms: 1) an ethical consent form which text is recommended and approved by the IRB, and 2) an additional legal consent form drafted by our institution's legal office that specifically addresses data usage, storage, and potential research applications.
- **Data anonymization.** All collected data were handled in accordance with the IRB requirements and privacy regulations at the authors' region. Personally identifiable information was removed before storage, and each subject was assigned a randomized subject identifier. No private attributes (*e.g.*, name, contact information) were retained in the dataset. Raw data were stored on secure, access-controlled servers, and only the de-identified iris images were used for model training and analysis. These measures ensure that the dataset cannot be linked back to individual subjects and that the privacy risks associated with biometric data are minimized.

Potential societal impact. Iris patterns constitute highly sensitive biometric identifiers, and the release of any iris dataset raises legitimate concerns regarding potential misuse, such as presentation attacks or unauthorized identity impersonation. We acknowledge these risks and therefore follow established practices in the biometric research community by adopting a controlled-access release of the dataset, further see Sec. H. Furthermore, the practical threat of impersonation using static iris images is limited, as modern iris recognition systems routinely incorporate liveness detection and anti-spoofing, making attacks based solely on still images significantly less feasible. All participants were informed of these risks and of our data protection measures during the consent process.

H. Dataset Availability

The dataset will be released very soon to promote the research viability of immersive iris recognition. It will be distributed under a controlled, non-commercial access protocol to assure ethical and legal compliance. Following the established practices in the biometric research community [4, 5, 17, 20, 23] research groups requesting the data are expected to provide institutional credentials and sign a legally binding agreement that restricts usage to academic research, prohibits any attempt at re-identification or misuse, and disallows redistribution to third parties. Access will be granted on a case-by-case basis to ensure compliance with these requirements.

References

[1] Worldcoin AI. Iris: Iris recognition inference system of the worldcoin project, 2023. 4

[2] Sunpreet S Arora, Mayank Vatsa, Richa Singh, and Anil Jain. On iris camera interoperability. In *2012 IEEE Fifth International Conference on Biometrics: Theory, Applications and Systems (BTAS)*, pages 346–352. IEEE, 2012. 2

[3] Sarah E Baker, Kevin W Bowyer, Patrick J Flynn, and P Jonathon Phillips. Template aging in iris biometrics. In *Handbook of Iris Recognition*, pages 205–218. Springer, 2012. 1

[4] CASIA. Casia iris image database version 1.0, 2002. 2, 1, 5

[5] CASIA. Casia iris database v4, 2018. 1, 2, 7, 5

[6] Chun-Nam Chun and Ronald Chung. Iris recognition for palm-top application. In *International Conference on Biometric Authentication*, pages 426–433. Springer, 2004. 1, 2

[7] John Daugman. How iris recognition works. In *The essential guide to image processing*, pages 715–739. Elsevier, 2009. 2, 3, 4, 5, 6, 7, 8

[8] Mihal Dobeš, Libor Machala, Petr Tichavskỳ, and J Pospíšil. Human eye iris recognition using the mutual information. *Optik*, 115(9):399–404, 2004. 1

[9] Melissa Edwards, Agnes Gozdzik, Kendra Ross, Jon Miles, and Esteban J Parra. Quantitative measures of iris color using high resolution photographs. *American journal of physical anthropology*, 147(1):141–149, 2012. 1

[10] Peter A Johnson, Paulo Lopez-Meyer, Nadezhda Sazonova, Fang Hua, and S Schuckers. Quality in face and iris research ensemble (q-fire). In *2010 Fourth IEEE International Conference on Biometrics: Theory, Applications and Systems (BTAS)*, pages 1–6. IEEE, 2010. 1

[11] Nassima Kihal, Salim Chitroub, Arnaud Polette, Isabelle Brunette, and Jean Meunier. Efficient multimodal ocular biometric system for person authentication based on iris texture and corneal shape. *IET Biometrics*, 6(6):379–386, 2017. 1

[12] Ajay Kumar and Arun Passi. Comparison and combination of iris matchers for reliable personal authentication. *Pattern recognition*, 43(3):1016–1026, 2010. 2, 1

[13] Kien Nguyen, Clinton Fookes, Sridha Sridharan, and Arun Ross. Complex-valued iris recognition network. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(1):182–196, 2022. 2, 3, 5, 6, 7, 8, 4

[14] Lubos Omelina, Jozef Goga, Jarmila Pavlovicova, Milos Oravec, and Bart Jansen. A survey of iris datasets. *Image and Vision Computing*, 108:104109, 2021. 1, 2, 4

[15] P Jonathon Phillips, Patrick J Flynn, Kevin W Bowyer, Richard W Vorder Bruegge, Patrick J Grother, George W Quinn, and Matthew Pruitt. Distinguishing identical twins by face recognition. In *2011 IEEE International Conference on Automatic Face & Gesture Recognition (FG)*, pages 185–192. IEEE, 2011. 1

[16] Hugo Proença and Luís A Alexandre. Ubiris: A noisy iris image database. In *International conference on image analysis and processing*, pages 970–977. Springer, 2005. 2, 1

[17] Hugo Proença, Silvio Filipe, Ricardo Santos, Joao Oliveira, and Luis A Alexandre. The ubiris. v2: A database of visible wavelength iris images captured on-the-move and at-a-distance. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(8):1529–1535, 2009. 2, 5, 1

[18] Ajita Rattani, Reza Derakhshani, Sashi K Saripalle, and Vikas Gottemukkula. Icip 2016 competition on mobile ocular biometric recognition. In *2016 IEEE international conference on image processing (ICIP)*, pages 320–324. IEEE, 2016. 2, 1

[19] Zhenan Sun and Tieniu Tan. Ordinal measures for iris recognition. *IEEE Transactions on pattern analysis and machine intelligence*, 31(12):2211–2226, 2008. 3, 5, 6, 7, 8, 4

[20] Kuo Wang and Ajay Kumar. Human identification in metaverse using egocentric iris recognition. *Authorea Preprints*, 2022. 1, 2, 3, 8, 4, 5

[21] Jianze Wei, Huaibo Huang, Yunlong Wang, Ran He, and Zhenan Sun. Towards more discriminative and robust iris recognition by learning uncertain factors. *IEEE Transactions on Information Forensics and Security*, 17:865–879, 2022. 2, 3, 5, 6, 7, 8, 4

[22] Jianze Wei, Yunlong Wang, Huaibo Huang, Ran He, Zhenan Sun, and Xingyu Gao. Contextual measures for iris recognition. *IEEE Transactions on Information Forensics and Security*, 18:57–70, 2022. 2, 3, 5, 6, 7, 8, 4

[23] Lihu Xiao, Zhenan Sun, Ran He, and Tieniu Tan. Coupled feature selection for cross-sensor iris recognition. In *2013 IEEE Sixth International Conference on Biometrics: Theory, Applications and Systems (BTAS)*, pages 1–6. IEEE, 2013. 2, 1, 5

[24] Yilong Yin, Lili Liu, and Xiwei Sun. Sdumla-hmt: A multimodal biometric database. In *Chinese conference on biometric recognition*, pages 260–268. Springer, 2011. 1

[25] Man Zhang, Qi Zhang, Zhenan Sun, Shujuan Zhou, and Nasir Uddin Ahmed. The btas competition on mobile iris recognition. In *2016 IEEE 8th international conference on biometrics theory, applications and systems (BTAS)*, pages 1–7. IEEE, 2016. 2, 1

[26] Qi Zhang, Haiqing Li, Zhenan Sun, and Tieniu Tan. Deep feature fusion for iris and periocular biometrics on mobile devices. *IEEE Transactions on Information Forensics and Security*, 13(11):2897–2912, 2018. 2, 3, 5, 6, 7, 8, 1, 4