

TEyeD: Over 20 million real-world eye images with Pupil, Eyelid, and Iris 2D and 3D Segmentations, 2D and 3D Landmarks, 3D Eyeball, Gaze Vector, and Eye Movement Types

WOLFGANG FUHL, University Tübingen, Germany
 GJERGJI KASNECI, University Tübingen, Germany
 ENKELEJDA KASNECI, University Tübingen, Germany

We present TEyeD, the world’s largest unified public data set of eye images taken with head-mounted devices. TEyeD was acquired with seven different head-mounted eye trackers. Among them, two eye trackers were integrated into virtual reality (VR) or augmented reality (AR) devices. The images in TEyeD were obtained from various tasks, including car rides, simulator rides, outdoor sports activities, and daily indoor activities. The data set includes 2D&3D landmarks, semantic segmentation, 3D eyeball annotation and the gaze vector and eye movement types for all images. Landmarks and semantic segmentation are provided for the pupil, iris and eyelids. Video lengths vary from a few minutes to several hours. With more than 20 million carefully annotated images, TEyeD provides a unique, coherent resource and a valuable foundation for advancing research in the field of computer vision, eye tracking and gaze estimation in modern VR and AR applications. Data and code at https://unitc-my.sharepoint.com/:f:/g/personal/iitfu01_cloud_uni-tuebingen_de/EvrNPdtigFVHtCMeFKSyLIUBepOcbX0nEkamweeZa0s9SQ?e=fWEvPp.

1 INTRODUCTION

Image-based eye tracking is becoming increasingly important in today’s world, as human eye movements have the potential to revolutionize the way we interact with computer systems around us [21, 66]. Since our actions and intentions can be recognized and - to a certain degree - anticipated from the way we move our eyes, eye movement analysis can enable completely new applications, especially when coupled with modern display technologies like VR or AR. For example, the gaze signal, together with the associated possibility of human-machine interaction [22], enables people with disabilities to interact with their environment through the use of special devices tailored to the patient’s disability [1]. In the case of surgical microscopes where the surgeon has to operate a multitude of controls, the visual signal can be used for automatic focusing [55, 60]. Furthermore, in scenarios where it is important to identify the expertise of a person (e.g., surgery, image interpretation, etc.) the gaze signal can be used together with interaction patterns to predict the expertise of a subject in a given task [2, 17, 18]. Gaze behavior can also be used to diagnose a variety of diseases [88], such as schizophrenia [65], autism [9], Alzheimer’s disease [88], glaucoma [68], and many more. Additionally, in VR/AR and gaming, the gaze signal can be used to reduce the computations of rendering resources [76].

A look at the human eye beyond the gaze information opens up further sources of information. For example, The gaze signal alone, however, is by no means the limit to information offered by the human eye [23, 27, 43, 45, 47]. The the frequency of eyelid closure [52–54], can be used to measure a person’s fatigue [15], an effective safety feature in driving [15] and aviation [14] scenarios. Of course, this applies to all safety critical tasks that are monitored by one or more persons [74]. Another significant source of information is the pupil size, which can serve as a basis to estimate the cognitive load of a person in a given task [16]. This may then be used for better adaption of the content (e.g. in media-based learning) to a person’s mental state. Finally, eye-related information can

be used in identification processes [23, 27], not only through the unique imprint of the iris [4], but also through an individual's gaze behavior [23, 24, 27, 51].

In the age of machine learning, where there is an abundance of effective and scalable learning approaches [6, 7, 19, 20, 73, 84, 87], it is, in principle, easier to develop algorithms or models which automatically retrieve the necessary information directly from the data. However, carefully annotated and curated data remains the central prerequisite for the development of machine learning - and especially deep learning - methods [37–39, 42] as well as the validation of the results [36, 49].

Providing such a prerequisite for a broad range of scenarios that involve eye-related information is exactly what the TEyeD data set aims to achieve.

Our contribution to the state-of-the-art is as follows:

- 1 We provide the largest unified data set of over 20 million eye images, collected using seven different eye trackers, ranging from 25 Hz to 200 Hz, including VR and AR.
- 2 Moreover, TEyeD covers a wide range of tasks and activities, such as car driving, driving in a simulator, indoor and outdoor activities.
- 3 Also provided is the ground truth of 3D landmarks and 3D segmentations not previously available for eye images..
- 4 TEyeD was generated from recordings in real-world settings, thus containing a wide range of realistic challenges, such as different resolutions, steep viewing angles, varying lighting conditions or device slippage.
- 5 TEyeD consists of both new images and image data from existing datasets. For the existing datasets, the link, citation, annotations and converter scripts are provided. Of the more than 20 million images, more than 15 millions are unpublished images, which we are allowed to provide.

2 RELATED WORK

Various eye-image data sets generated by eye trackers already exist [46, 56–58, 63, 70, 89, 93] including recordings from driver studies as well as simulators [46, 56–58]. In addition, there are also recordings from specific challenges [89, 93] as well as real time capable algorithms [29, 31, 33, 34, 44?]. While early data sets provided only the annotated pupil center [46, 56–58, 61, 63, 70, 89, 93], newer data sets offer the segmented pupil, iris, and sclera [32, 50, 63], eventually extended by the optical vector [30, 55] which allows for a shift in invariant gaze estimation [90]. Such data sets are available for conventional eye trackers [46, 56–58, 63, 70, 89, 93] and for VR [63, 70] and AR [70]. In addition to these annotations, segmented iris data sets with subject identification numbers are available for the development of personal identification systems [77–79]. Other annotations that contain important information are eye movement types such as fixation, saccades, and smooth pursuits [72]. However, in contrast to TEyeD, all the mentioned data sets have a narrow, task-specific focus.

Since the manual annotation of eye images and eye movements is very complex, especially when a high accuracy is required, several procedures to generate synthetic data have been proposed [70, 91, 95]. This includes synthesized image data [70], automated rendering methods [91, 95], and eye movement simulations [35, 48, 59], as well as generative adversarial networks (GAN) [32]. The disadvantage of synthetic data sets is that they cannot represent relevant challenges of real-world imaging, e.g., with regard to varying illumination conditions, physiological properties of the eyes, lighting sources, device slippage, and more which makes it difficult for the algorithms to detect eye movements [25, 26, 40]. This is still an important part of research today, particularly in the development of novel interaction techniques and applications in VR/AR.

In summary, current published data sets are limited due to their focus on specific problems. To the best of our knowledge, there is no unified and coherent data set containing all the relevant

annotations on eye-related information. Moreover, all data sets are generated using one specific type of eye-tracking device. In contrast, TEyeD offers a carefully, coherently and richly labeled data set containing all relevant eye-related information on a wide range of tasks (such as car driving, driving in a simulator, indoor and outdoor activities including VR and AR scenarios). The tasks, in total, were recorded by seven different eye trackers with different recording frequencies (i.e. sampling rates ranging from 25 to 200 Hz): One for VR, one for AR, and five more from head-mounted devices. In addition, TEyeD was generated from recordings containing a wide range of realistic challenges, such as different resolutions, steep viewing angles and varying lighting conditions or device slippage during outdoor and sports activities, to name a few. These challenges are known to be the limiting factor of eye-tracking in many real-world applications [62].

3 COMPARISON WITH EXISTING DATA SETS

Table 1 provides an overview of existing data sets containing close-up eye images. Each data set deals with a specific issue, such as Casia and Ubiiris, which are used to identify individuals by the iris. Direct gaze estimation, as in POG, NVGaze or NNVEC, is tackled by a more recent group of data sets. In NNVEC, the direct estimation of the optical vector and eyeball position make it possible to compensate for shifts of the head mounted eye tracker. In contrast to Casia and Ubiiris, MASD focuses on the segmentation of the eye’s sclera. MASD can be used to improve iris segmentation while also helping to estimate the degree of eye opening, an indicator of blink rate. Different eye movement types are offered alongside images in GIW and BAY while additional annotations of eye movement types for worn eye trackers are published in HEV and HEI. Proving to be very challenging, GAN and 550k existing data sets, LPW, ExCuSe, Else, and PNET, were extended with segmentations for the pupil and sclera. Originally, these data sets, together with the Swi data set, only provided the pupil center as annotation. OpenEDS was the first data set with segmentations for the pupil, iris, and sclera. Containing many subjects, OpenEDS was specifically acquired to enable VR-related research and applications. Additionally, two data sets (EWO and FRE) encompass 2D landmarks specific to worn eye trackers and were published together with real-time algorithms for the CPU.

TEyeD both combines and extends previously published data sets by utilizing seven different eye trackers, each with a different resolution, incorporating all available annotations offered by existing data sets, and broadening these sets with 3D segmentation and landmarks. More specifically, the data sets integrated in TEyeD are NNGaze, LPW, GIW, ElSe, ExCuSe, and PNET. Additionally, the complete data from the study [69] was also carefully annotated. Additional annotated data was recorded with an eye tracker from Enke GmbH and the eye tracker from Look!. In total, TEyeD contains more than 20 million images, making it, to our knowledge, the world’s largest data set of images taken by head mounted eye trackers. Similar to the OpenEDS data set, some data (6,379,400 samples collected from the Look! and Enke GmbH eye trackers) was withheld from TEyeD in order to obtain a reliable evaluation beyond generalization of single eye trackers. By deploying trained models or pre-compiled programs, it is possible to achieve fair evaluation of the runtime and generalization of different eye trackers. Unfortunately, we could not consider eye images captured with Tobii eye trackers due to prohibitory license agreements.

4 FURTHER DATA SET DETAILS

Figure 1 shows sample images of TEyeD. The first and fifth column contain the input images. The second and sixth column show these images with overlaid segmentations of the sclera, iris and pupil. The third and seventh column show the landmarks on the input image with red landmarks belonging to the eyelids, green landmarks to the iris, and white landmarks to the pupil. In the fourth and eighth column, the calculated eyeball is displayed as well as the center of the eyeball and the gaze vector.

Table 1. A list of the published data sets for virtual reality (VR), augmented reality (AR), and head mounted (HM) eye tracker. The table contains information on the number of subjects (Sub.), the type of eye tracker (AR,VR,HM), the acquisition frequency (FRQ), image resolution (Res.), the number of annotated images (Num. Annot), whether or not segmentations are present in 2D&3D (Seg 2D, Seg 3D), whether or not the pupil center is annotated (PC), whether or not landmarks are present in 2D&3D (LM 2D,LM 3D), whether the position and radius of the eyeball is given (Eye), whether or not the gaze vector or gaze position is given (Ga), and whether or not the eye movement types are annotated (Mov.). The subtypes stand for *I* = *Iris*, *P* = *Pupil*, *Sc* = *Sclera*, *Lid* = *Eyelid*, *F* = *Fixation*, *S* = *Saccade*, *SP* = *Smooth Pursuits*, and *B* = *Blinks*.

Data	Sub.	Tracker VR AR HM			FRQ	Res.	Num. Annot	Seg 2D P I Sc			Seg 3D P I Sc			PC	LM 2D P I Lid			LM 3D P I Lid			Eye	Ga	Mov. F S SP B				
POG[75]	20	-	-	1	30Hz	768 × 480	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	Y	-	-	-	-	-
NNVEC[30]	20	-	-	1	25Hz	384 × 288	866,069	-	-	-	-	-	-	-	-	-	-	-	-	-	Y	Y	-	-	-	-	-
NVGaze[70]	35	1	1	-	120Hz	640 × 480	2,500,000	-	-	-	-	-	-	-	-	-	-	-	-	-	Y	Y	-	-	-	-	-
Casia.v1[77, 92]	108	-	-	1	-	320 × 280	756	-	Y	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
Casia.v2[77, 92]	60	-	-	2	-	640 × 480	2,400	-	Y	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
Casia.v3[77, 92]	≈700	-	-	3	-	Multiple	22,034	-	Y	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
Casia.v4[77, 92]	≈1,800	-	-	4	-	Multiple	54,601	-	Y	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
Casia.test[77, 92]	1,000	-	-	1	-	640 × 480	10,000	-	Y	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
Casia.age[3, 94]	50	-	-	2	-	Multiple	≈160,000	-	Y	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
Ubiris.v1[78]	241	-	-	Y	30Hz	Multiple	877	-	Y	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
Ubiris.v2[79]	261	-	-	Y	200Hz	400 × 300	≈11,000	-	Y	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
MASD[10–12]	82	-	-	Y	-	Multiple	2,624	-	-	Y	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
GAN[32]	22	-	-	1	120Hz	640 × 480	130,856	Y	-	Y	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
500k[50]	20	-	-	1	25Hz	384 × 288	866,069	Y	-	Y	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
MEMD[41]	20	-	-	1	25Hz	384 × 288	866,069	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	Y	Y	Y	-
ME[96]	587	1	1	-	-	640 × 480	880,000	Y	Y	Y	-	-	-	-	-	-	-	-	-	-	-	Y	-	-	-	-	-
OpenEDS[63]	152	1	-	-	200Hz	400 × 640	356,649	Y	Y	Y	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
GIW[72]	19	-	-	1	120Hz	640 × 480	≈2,016,000	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	Y	Y	Y	Y	Y
BAY[86]	6	-	-	1	30Hz	640 × 480	27,022	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	Y	Y	Y	-	-
HEV[13]	57	1	-	-	24-30Hz	-	no images	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	Y	Y	-	-	-
HEI[82]	63	1	-	-	60Hz	-	no images	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	Y	Y	-	-	-
LPW[93]	22	-	-	1	120Hz	640 × 480	130,856	-	-	-	-	-	-	Y	-	-	-	-	-	-	-	-	-	-	-	-	-
Swi[89]	2	-	-	1	-	620 × 460	600	-	-	-	-	-	-	Y	-	-	-	-	-	-	-	-	-	-	-	-	-
ExCuSe[46]	7	-	-	1	25Hz	384 × 288	39,001	-	-	-	-	-	-	Y	-	-	-	-	-	-	-	-	-	-	-	-	-
Else[58]	17	-	-	1	25Hz	384 × 288	55,712	-	-	-	-	-	-	Y	-	-	-	-	-	-	-	-	-	-	-	-	-
PNET[56, 57]	5	-	-	1	25Hz	384 × 288	41,217	-	-	-	-	-	-	Y	-	-	-	-	-	-	-	-	-	-	-	-	-
EWO[53]	11	-	-	1	25Hz	384 × 288	1,100	-	-	-	-	-	-	-	-	Y	-	-	-	-	-	-	-	-	-	-	-
FRE[54]	11	-	-	1	25Hz	384 × 288	4,000	-	-	-	-	-	-	-	-	Y	-	-	-	-	-	-	-	-	-	-	-
TEyeD	39	-	-	1	25Hz	384 × 288	5,665,053	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
	1	-	-	1	60Hz	320 × 240	12,184	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
	22	-	-	1	95Hz	640 × 480	130,856	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
	54	1	1	1	120Hz	640 × 480	8,691,764	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
	16	-	-	1	60Hz	640 × 360	6,367,216	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y

Table 2 shows rough statistics from the TEyeD data set. Interestingly, our data set contains also images in which no eye is present. This can occur when the eye tracker is removed from the subject or when reflections in the near-infrared rage of the subject’s glasses are so strong that the eye is no longer visible. Finally, our data set contains images in which the pupil is annotated but no iris appears. Such images can occur when the eye tracker is removed from the subject or when reflections in the near-infrared rage of the subject’s glasses are so strong that the eye is no longer visible.




Fig. 1. Example images from our data set with annotations.

Figures 2 shows the logarithmic distribution of landmarks for the pupil (left), iris (middle), and eyelids (right). Since a large part of our image data comes from real images, we used the logarithm to show all occurrences, since this enables a better representation of the areas, which are underrepresented in normal gaze behavior. One such occurrence, for example, could be a subject




Fig. 2. The logarithmic distribution of the pupil landmarks (left), iris landmarks (middle), and eyelid landmarks (right) in TEyeD.

Table 2. General statistics of our data set.

Pupils	19,927,927
Iris	19,756,546
Eyelid	20,666,096
No eye images	200,977
Open eyes	19,859,456
Closed eyes	806,640

driving a car. The logarithm accounts for the driver’s view mainly being directed forward. Hereby, our data set can also be used to evaluate tracking algorithms. In addition to the logarithmic distribution, black crosses represent the mean position of all landmarks in Figures 2, which are distributed over almost the entire image area. There are also individual landmarks located outside of the image area, especially in the corners of the eyelids and in the upper landmarks of the iris, as can be seen in Figure 2.

Figure 3 shows the area distribution (in pixels) of the pupil, iris, and the sclera as a whisker plot. The blue boxes represent the confidence intervals with the 25th and 75th percentiles. In the middle of the blue box, the red line represents the median. The red crosses represent. As exhibited here, our data set contains different camera distances due to the specificities of the individual eye trackers and the LPW data set’s special recordings which were taken at close range. TEyeD also incorporates both large and small pupils, the result of different camera distances as well as a variety of lighting conditions.

Figure 3 shows the logarithmic gaze vector distribution, where all vectors are unit vectors and shifted to the same center. As this figure incorporates close-up images of the eye, the depth of every vector favors the direction of the camera. Thus, depth information is not shown separately. We decided to use a logarithmic representation, similar to the landmarks, because the gaze is typically consistent and centrally aligned during activities such as driving a car. This also allows for the evaluation of tracking algorithms on TEyeD. As shown in Figure 3, the gaze vector is distributed over the entirety of the eye ball hemisphere.

Figure 4 shows the eyeball position (x,y) distribution as well as the eyeball radius in pixels as a whisker plot mapped to a fixed resolution of 192×144 . The z position is not shown because we set the z position for the eye balls to zero. This means that the eyeball center is the origin for the 3D positions of the landmarks and 3D segmentations. As exhibited in Figure 4, most of the eyeball centers are located in the image. There are, however, some eyeball centers located outside of the image (Y position greater than 144). As is the case for some of the images in the LPW data set, this




Fig. 3. Left: The area distribution for the pupil, iris, and eyelids on an 192×144 image resolution. The blue box corresponds to the 25th and 75th percentiles. Red crosses are the outlier and the red line corresponds to the median.

Right: The logarithmic distribution of the gaze vector centered and mapped to a unit sphere in our data set.




Fig. 4. The distributions for the eyeball center x and y as well as the distribution of the eyeball radius in our data set on a 192×144 image resolution. The blue box corresponds to the 25th and 75th percentiles. Red crosses are the outlier and the red line corresponds to the median.

is due to the camera’s very close proximity to the eye. The wide variation in the eyeball’s radius is likewise a result of the camera’s different distances.

5 ANNOTATION PROCESS

For the annotation of the landmarks and semantic segmentation in TEyeD, we used a semi-supervised approach together with the multiple annotation maturation (MAM) [28] algorithm. Unlike the original algorithm, we used CNNs [73, 84] instead of SVMs [87] in combination with HOG [8] features. We also limited the iterations to five and used two competing models. One model consisted of a ResNet50 and was trained for landmark regression using the validation loss function of [36]. This loss function enabled the CNN to detect whether the pupil, iris, and eyelids were present. The loss function also provided information about the accuracy of the individual landmarks. For the other model, we trained the semantic segmentation together with a U-Net [83] and residual blocks [64]. For both models, we also used the batch balancing of [36].

Initially, we annotated 20,000 images with landmarks and converted them into semantic segmentations. Then we trained the CNNs and continuously improved them with the MAM algorithm. After five iterations, the ResNet50 landmarks were converted into semantic segmentations and compared to the U-Net results. For this step, we used the Jaccard index, i.e., $\frac{\text{ResNet50} \cap \text{U-Net}}{\text{ResNet50} \cup \text{U-Net}}$. If this value was less than 0.9, applicable images were marked and new images selected from the set for manual annotation. A total of four post-annotations were completed and the process was started again from scratch.

3D eyeball and optical vector were annotated based on the approach presented in [30]. However, instead of using the pupil ellipse, we used the iris ellipse since it is only partially affected by corneal refraction. Additionally, we used both approaches from [30], wherein one approach processes several ellipses in a neural network and the other approach calculates single vectors from single ellipses. For the second approach, we calculated the minimum intersection point of the individual vectors and the resulting radius. In regard to segmentation, we compared both approaches and accepted deviations of less than two pixels for the center and radius of the eyeball. In all other cases, we made manual correction utilizing the preceding and succeeding eyeball parameters.

3D landmarks and segmentation were calculated geometrically by combining the 2D landmarks and segmentations with the 3D eyeball model. As the pupil is always physically located at the center of the iris, we accounted for two different 3D segmentations and 3D landmarks. We first considered how the pupil appears in the eye image. Due to corneal refractions and steep camera angles, the pupil appears often not in the center of the iris. Accordingly, we adjusted the 3D landmarks and 3D segmentation to the iris and, more specifically, to the center of the iris.

Eye movements are annotated as fixations ("still" eye), saccades (fast eye movement between fixations), smooth pursuits (slow eye movement), and blinks. Additionally, all images without an eye or with open eyes lacking valid pupil coverage were marked as errors. In the first step of our annotation, 50,000 individual images were annotated using the optical vector annotation. Then, semantic segmentation for eye movements was applied to the angular velocities of the optical vector [48]. On top of it, we applied the MAM approach for two iterations. Finally, the detected eye movement types were validated against biologically valid parameters [80] and manually corrected for errors.

6 BASELINE EVALUATIONS

Generalisation across eye trackers. To highlight some of the advantages that come with this large data set, in our first baseline experiment we analyzed the generalization performance for **landmark regression** and **semantic segmentation** across different eye trackers. Note that cross-eye-tracker generalization poses a key challenge for eye-tracking manufacturers for the mentioned tasks, since, as of now, changing eye-tracking devices involves the manual annotation of images generated by the new device.

In our experiment, we split the data into a validation and training set. This is done by assigning 50% of the recordings into each set, excluding the hold back images (6,379,400 from the eye trackers Look! and Enke GmbH) which are used for the final evaluation. In order to avoid the same subjects in the training and validation set, whole recordings were always assigned to either the training or the validation set.

As a test data set, we hold back 6,379,400 images with annotations from the eye trackers Look! and Enke GmbH. In order to evaluate over this data, we used the models ResNet-34 [64], ResNet-50 [64], MobilNetV2 [85], and U-Net [83] with residual blocks [64] and batch normalization [67]. The training data was additionally augmented with 0-30% random noise, rotations between -45° - 45° , shifts of 0-20%, 1.0-2.0 standard deviation blure, overlaying with images to simulate reflections, adding vertical and horizontal noise to pixel lines, and adding 0-10 noisy squares or ellipses with random size and orientation. As optimizer for the semantic segmentation, we used SGD [81] with

Table 3. Landmark regression results in average euclidean pixel distance divided by the image resolution diagonal and multiplied with the factor 10^2 for the pupil, iris, and eyelid 3D landmarks on TEyeD. Best results in bold.

Model	Train Data	Res.	Pupil	Iris	Eyelid
		px	$\frac{px}{diag}$	$\frac{px}{diag}$	$\frac{px}{diag}$
Res-34	LPW	192×144	3.23	2.70	2.51
		384×288	2.92	2.48	2.19
		768×576	2.45	2.07	1.83
Res-34	ALL	192×144	2.15	1.85	1.55
		384×288	1.89	1.63	1.25
		768×576	1.76	1.46	1.13
Res-50	ALL	192×144	2.02	1.65	1.34
		384×288	1.68	1.46	1.08
		768×576	1.54	1.34	1.01
MobV2	ALL	192×144	2.50	2.01	1.62
		384×288	2.11	1.73	1.40
		768×576	1.94	1.54	1.20

the parameters $5 * 10^{-4}$ weight decay, 0.99 momentum, and 0.1 learning rate. After every sequence one thousand epochs, the learning rate was reduced by a factor of 0.1. This was performed up to a learning rate of 10^{-5} . As a loss function for the pixel classes softmax was used. For landmark regression, Adam [71] was used with the parameters $5 * 10^{-4}$ as weight decay, 0.9 and 0.99 for the first and second momentum, respectively, and learning rate of 10^{-4} . After every sequence of one thousand epochs, the learning rate was reduced by a factor of 0.1. This was enacted up to a learning rate of 10^{-8} . L2 was used as the loss function.

Evaluation environment We used the C++-based CuDNN framework for the neural net models. The hardware for the test environment involves an Intel i5-4570 CPU with 4 cores, 16 Gigabyte DDR4 memory and an NVIDIA 1050ti with 4 Gigabyte memory.

Results on landmark regression. Table 3 shows the results of the landmark regression. For this purpose, we trained different models that determine landmarks for the pupil, iris, and eyelids together. Note, however, that the results can be further improved by using individual models for estimating the landmarks of the pupil, iris, and the eyelids. This is largely because eyelids move independently of the pupil and the iris, and the pupil is displaced from the iris due to corneal refraction. As an evaluation measure we report the mean distance of the predictions from the ground truth annotations, as pixels normalized by the diagonal (as $\frac{px}{diag}$). Table 3 shows, as expected, that larger models are more effective on the described regression tasks. The same conclusion can be drawn from Table 4, where the results of the eyeball parameter estimation are shown. For this purpose, we trained different models, each having received five consecutive images as input. Also in this case, larger models and higher resolutions are more effective. However, in both Tables 3 and 4 we can see the clear advantage of the TEyeD data set in comparison to smaller existing data sets, as depicted by the topmost two models (highlighted in gray) which use the same model architecture, i.e., ResNet-34, but are trained once on the LPW data set and once the full TEyeD data set. Furthermore, the results also indicate, as expected, that cross-eye-tracker generalization on images taken in real-world settings is a challenging task, which however can be approached using TEyeD together with more complex architectures. Thus, now the key challenge of cross-eye-tracker generalization can be easily approached without the need for creating and annotating new data, whenever a new eye-tracking device is used.

Table 4. Eyeball parameter and gaze vector (GV) regression results in average euclidean pixel distance divided by the image resolution diagonal and multiplied with the factor 10^2 for the 3D position as well as for the radius and average angular difference for the gaze vector on TEyeD. Each model received five consecutive images as input to estimate the eye ball parameters and the current gaze vector. Best results in bold.

Model	Train Data	Res.	Eyeball	Radius	GV
		px	$\frac{px}{diag}$	$\frac{px}{diag}$	degree
Res-34	LPW	192×144	2.41	2.19	5.37
		384×288	2.16	2.02	4.90
		768×576	1.80	1.93	4.41
Res-34	ALL	192×144	1.74	1.35	4.72
		384×288	1.54	1.22	3.92
		768×576	1.32	1.03	3.18
Res-50	ALL	192×144	1.54	1.22	4.15
		384×288	1.25	1.06	3.69
		768×576	1.05	0.96	2.97
MobV2	ALL	192×144	1.95	1.72	5.05
		384×288	1.73	1.62	4.19
		768×576	1.55	1.25	3.57

Table 5. Semantic segmentation results as mean Jaccard index (mJI) on the test set. For the models ResNet-34 (Res-34), ResNet-50 (Res-50), and MbiNetV2 (MobV2) we converted the landmarks into segments using OpenCV [5]. Z is the average euclidean distance divided by the image resolution diagonal and multiplied with the factor 10^2 for the 3D position of the segments. Best results in bold.

Model	Res.	Pupil	Iris	Sclera	Z
	px	mJI	mJI	mJI	$\frac{px}{diag}$
U-Net	192×144	0.52	0.58	0.67	2.73
	384×288	0.56	0.61	0.69	2.49
	768×576	0.60	0.62	0.70	1.95
Res-34	192×144	0.58	0.62	0.71	1.90
	384×288	0.60	0.64	0.73	1.45
	768×576	0.63	0.68	0.76	1.12
Res-50	192×144	0.61	0.65	0.75	1.71
	384×288	0.64	0.66	0.77	1.22
	768×576	0.65	0.70	0.78	0.85
MobV2	192×144	0.56	0.60	0.70	2.07
	384×288	0.59	0.62	0.72	1.75
	768×576	0.61	0.65	0.74	1.46

Semantic segmentation. Table 5 shows the results for semantic segmentation. For the landmark regression models, we created the semantic segments using the OpenCV ellipse fit for the iris and the pupil, and the fillPolygon function for the eyelids. Also for this task, we conclude that despite the challenging eye images from different eye trackers and real-world scenarios, a fairly viable generalization can be achieved using TEyeD together with larger models.

Table 6. Eye movement segmentation results are provided as the mean Jaccard Index (mJI) on the test set. The models ResNet-34, ResNet-50, and MobileNetV2 are used in a window-based fashion on 256 consecutive input values from the ground truth, i.e., pupil center (PC) or gaze vector (GV), and predict on 16 consecutive data points, i.e. the corresponding eye movement events: Fixations (Fix.), Saccades (Sacc.), Smooth Pursuits (Sm.Purs.), Errors and Blinks. Best results are highlighted in bold.

Input	Model	Fix.	Sacc.	Sm.Purs.	Error	Blink
PC	ResNet-34	0.81	0.73	0.83	0.92	0.81
	ResNet-50	0.86	0.75	0.87	0.95	0.83
	MobileNetV2	0.78	0.70	0.81	0.89	0.74
GV	ResNet-34	0.92	0.87	0.91	0.98	0.90
	ResNet-50	0.94	0.89	0.93	0.98	0.91
	MobileNetV2	0.85	0.82	0.89	0.97	0.88

Recognition of eye movement types. Table 6 presents the results on eye movement recognition. The models had to predict the eye movement type in addition to the errors and blinks. For this purpose, they received the ground truth of the pupil in the upper part of the evaluation and the gaze vector in the second part of the evaluation. All models were applied in a window-based fashion and received 256 data points (raw eye-tracking data points) to classify 16 data points. These 16 data points to be predicted were exactly in the middle of the 256 data points. For each data point, we also appended the time in milliseconds (ms) from the previous data point. As it can be seen, the gaze vector (GV) is much more effective for eye-movement classification because it compensates for shifts of the eye tracker. Due to the difficulty of computing a robust signal of the gaze vector, the pupil center is still taken in conventional systems. Also in this case, TEyeD can be used to achieve the generalization across different eye trackers and different real-world settings.

7 CONCLUSION

In this work, we presented TEyeD, a rich and coherent data set of over 20 Million eye images along with their 2D and 3D annotations and other annotations including eye movement types, semantic segmentations, landmarks, elliptical parameters for the iris, the pupil, and the eyelid as well as eyeball parameters for shift invariant gaze estimation. Generated by a total of seven different eye trackers with different sampling rates and under challenging real-world conditions, TEyeD is the most comprehensive and realistic data set of semantically annotated eye images to date. This data set should not only be seen as a new foundational resource in the field of computer vision and eye movement research. We are convinced that TEyeD will also have a profound impact on other fields and communities, ranging from cognitive sciences to AR and VR applications. At the very least, it will unquestionably contribute to the application of eye-movement and gaze estimation techniques in challenging practical use cases.

REFERENCES

- [1] Malek Adjouadi, Anaelis Sestin, Melvin Ayala, and Mercedes Cabrerizo. 2004. Remote eye gaze tracking system as a computer interface for persons with severe motor disability. In *International conference on computers for handicapped persons*. Springer, 761–769.
- [2] H. Bahmani, W. Fuhl, E. Gutierrez, G. Kasneci, E. Kasneci, and S. Wahl. 2016. Feature-based attentional influences on the accommodation response. In *Vision Sciences Society Annual Meeting Abstract*.
- [3] Thomas Bergmüller, Luca Debiasi, Andreas Uhl, and Zhenan Sun. 2014. Impact of sensor ageing on iris recognition. In *IEEE International Joint Conference on Biometrics*. IEEE, 1–8.

- [4] Wageeh W Boles. 1998. A security system based on human iris identification using wavelet transform. *Engineering Applications of Artificial Intelligence* 11, 1 (1998), 77–85.
- [5] G. Bradski. 2000. The OpenCV Library. *Dr. Dobbs's Journal of Software Tools* (2000).
- [6] Leo Breiman. 2001. Random forests. *Machine learning* 45, 1 (2001), 5–32.
- [7] Tianqi Chen and Carlos Guestrin. 2016. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*. 785–794.
- [8] Navneet Dalal and Bill Triggs. 2005. Histograms of oriented gradients for human detection. In *2005 IEEE computer society conference on computer vision and pattern recognition (CVPR'05)*, Vol. 1. IEEE, 886–893.
- [9] Kim M Dalton, Brendon M Nacewicz, Tom Johnstone, Hillary S Schaefer, Morton Ann Gernsbacher, Hill H Goldsmith, Andrew L Alexander, and Richard J Davidson. 2005. Gaze fixation and the neural circuitry of face processing in autism. *Nature neuroscience* 8, 4 (2005), 519–526.
- [10] Abhijit Das, Umapada Pal, Michael Blumenstein, Caiyong Wang, Yong He, Yuhao Zhu, and Zhenan Sun. 2019. Sclera segmentation benchmarking competition in cross-resolution environment. In *2019 International Conference on Biometrics (ICB)*. IEEE, 1–7.
- [11] Abhijit Das, Umapada Pal, Miguel A Ferrer, and Michael Blumenstein. 2016. SSRBC 2016: Sclera segmentation and recognition benchmarking competition. In *2016 International Conference on Biometrics (ICB)*. IEEE, 1–6.
- [12] Abhijit Das, Umapada Pal, Miguel A Ferrer, Michael Blumenstein, Dejan Štepec, Peter Rot, Žiga Emeršič, Peter Peer, Vitomir Štruc, SV Aruna Kumar, et al. 2017. SSERBC 2017: Sclera segmentation and eye recognition benchmarking competition. In *2017 IEEE International Joint Conference on Biometrics (IJCB)*. IEEE, 742–747.
- [13] Erwan J David, Jesús Gutiérrez, Antoine Coutrot, Matthieu Perreira Da Silva, and Patrick Le Callet. 2018. A dataset of head and eye movements for 360 videos. In *Proceedings of the 9th ACM Multimedia Systems Conference*. 432–437.
- [14] David F Dinges, Greg Maislin, Rebecca M Brewster, Gerald P Krueger, and Robert J Carroll. 2005. Pilot test of fatigue management technologies. *Transportation research record* 1922, 1 (2005), 175–182.
- [15] Wenhui Dong and Xiaojuan Wu. 2005. Fatigue detection based on the distance of eyelid. In *Proceedings of 2005 IEEE International Workshop on VLSI Design and Video Technology, 2005*. IEEE, 365–368.
- [16] Andrew T Duchowski, Krzysztof Krejtz, Izabela Krejtz, Cezary Biele, Anna Niedzielska, Peter Kiefer, Martin Raubal, and Ioannis Giannopoulos. 2018. The index of pupillary activity: Measuring cognitive load vis-à-vis task difficulty with pupil oscillation. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–13.
- [17] S Eivazi, W Fuhl, and E Kasneci. 2017. Towards intelligent surgical microscopes: Surgeons gaze and instrument tracking. In *Proceedings of the 22st International Conference on Intelligent User Interfaces, IUI*.
- [18] S. Eivazi, A. Hafez, W. Fuhl, H. Afkari, E. Kasneci, M. Lehecka, and R. Bednarik. 2017. Optimal eye movement strategies: a comparison of neurosurgeons gaze patterns when using a surgical microscope. *Acta Neurochirurgica* (2017).
- [19] Yoav Freund, Robert E Schapire, et al. 1996. Experiments with a new boosting algorithm. In *icml*, Vol. 96. Citeseer, 148–156.
- [20] Jerome H Friedman. 2002. Stochastic gradient boosting. *Computational statistics & data analysis* 38, 4 (2002), 367–378.
- [21] W. Fuhl. 2019. *Image-based extraction of eye features for robust eye tracking*. Ph.D. Dissertation. University of Tübingen.
- [22] Wolfgang Fuhl. 2020. From perception to action using observed actions to learn gestures. *User Modeling and User-Adapted Interaction* (08 2020), 1–18.
- [23] Wolfgang Fuhl, Efe Bozkir, Benedikt Hosp, Nora Castner, David Geisler, Thiago C Santini, and Enkelejda Kasneci. 2019. Encodji: encoding gaze data into emoji space for an amusing scanpath classification approach. In *Proceedings of the 11th ACM Symposium on Eye Tracking Research & Applications*. 1–4.
- [24] Wolfgang Fuhl, Efe Bozkir, and Enkelejda Kasneci. 2020. Reinforcement learning for the privacy preservation and manipulation of eye tracking data. *arXiv preprint arXiv:2002.06806* (08 2020).
- [25] W. Fuhl, N. Castner, and E. Kasneci. 2018. Histogram of oriented velocities for eye movement detection. In *International Conference on Multimodal Interaction Workshops, ICMIW*.
- [26] W. Fuhl, N. Castner, and E. Kasneci. 2018. Rule based learning for eye movement type detection. In *International Conference on Multimodal Interaction Workshops, ICMIW*.
- [27] W. Fuhl, N. Castner, T. C. Kübler, A. Lotz, W. Rosenstiel, and E. Kasneci. 2019. Ferns for area of interest free scanpath classification. In *Proceedings of the 2019 ACM Symposium on Eye Tracking Research & Applications (ETRA)*.
- [28] W. Fuhl, N. Castner, L. Zhuang, M. Holzer, W. Rosenstiel, and E. Kasneci. 2018. MAM: Transfer learning for fully automatic video annotation and specialized detector creation. In *International Conference on Computer Vision Workshops, ICCVW*.
- [29] W. Fuhl, S. Eivazi, B. Hosp, A. Eivazi, W. Rosenstiel, and E. Kasneci. 2018. BORE: Boosted-oriented edge optimization for robust, real time remote pupil center detection. In *Eye Tracking Research and Applications, ETRA*.

- [30] W. Fuhl, H. Gao, and E. Kasneci. 2020. Neural networks for optical vector and eye ball parameter estimation. In *ACM Symposium on Eye Tracking Research & Applications, ETRA 2020*. ACM.
- [31] W. Fuhl, H. Gao, and E. Kasneci. 2020. Tiny convolution, decision tree, and binary neuronal networks for robust and real time pupil outline estimation. In *ACM Symposium on Eye Tracking Research & Applications, ETRA 2020*. ACM.
- [32] W. Fuhl, D. Geisler, W. Rosenstiel, and E. Kasneci. 2019. The applicability of Cycle GANs for pupil and eyelid segmentation, data generation and image refinement. In *International Conference on Computer Vision Workshops, ICCVW*.
- [33] W. Fuhl, D. Geisler, T. Santini, T. Appel, W. Rosenstiel, and E. Kasneci. 2018. CBF: Circular binary features for robust and real-time pupil center detection. In *ACM Symposium on Eye Tracking Research & Applications*.
- [34] W. Fuhl, D. Geisler, T. Santini, and E. Kasneci. 2016. Evaluation of State-of-the-Art Pupil Detection Algorithms on Remote Eye Images. In *ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct publication – PETMEI 2016*.
- [35] W. Fuhl and E. Kasneci. 2018. Eye movement velocity and gaze data generator for evaluation, robustness testing and assess of eye tracking software and visualization tools. In *Poster at Egocentric Perception, Interaction and Computing, EPIC*.
- [36] W. Fuhl and E. Kasneci. 2019. Learning to validate the quality of detected landmarks. In *International Conference on Machine Vision, ICMV*.
- [37] Wolfgang Fuhl and Enkelejda Kasneci. 2020. Multi Layer Neural Networks as Replacement for Pooling Operations. *arXiv preprint arXiv:2006.06969* (08 2020).
- [38] Wolfgang Fuhl and Enkelejda Kasneci. 2020. Rotated Ring, Radial and Depth Wise Separable Radial Convolutions. *arXiv preprint arXiv:2010.00873* (08 2020).
- [39] Wolfgang Fuhl and Enkelejda Kasneci. 2020. Weight and Gradient Centralization in Deep Neural Networks. *arXiv preprint arXiv:2010.00866* (08 2020).
- [40] W Fuhl and E Kasneci. 2021. A Multimodal Eye Movement Dataset and a Multimodal Eye Movement Segmentation Analysis. *arXiv preprint arXiv:2101.04318* (01 2021).
- [41] Wolfgang Fuhl and Enkelejda Kasneci. 2021. A Multimodal Eye Movement Dataset and a Multimodal Eye Movement Segmentation Analysis. *arXiv preprint arXiv:2101.04318* (2021).
- [42] W. Fuhl, G. Kasneci, W. Rosenstiel, and E. Kasneci. 2020. Training Decision Trees as Replacement for Convolution Layers. In *Conference on Artificial Intelligence, AAI*.
- [43] W. Fuhl, T. C. Kübler, H. Brinkmann, R. Rosenberg, W. Rosenstiel, and E. Kasneci. 2018. Region of interest generation algorithms for eye tracking data. In *Third Workshop on Eye Tracking and Visualization (ETVIS), in conjunction with ACM ETRA*.
- [44] W. Fuhl, T. C. Kübler, D. Hospach, O. Bringmann, W. Rosenstiel, and E. Kasneci. 2017. Ways of improving the precision of eye tracking data: Controlling the influence of dirt and dust on pupil detection. *Journal of Eye Movement Research* 10, 3 (05 2017).
- [45] W. Fuhl, T. C. Kübler, K. Sippel, W. Rosenstiel, and E. Kasneci. 2015. Arbitrarily shaped areas of interest based on gaze density gradient. In *European Conference on Eye Movements, ECEM 2015*.
- [46] W. Fuhl, T. C. Kübler, K. Sippel, W. Rosenstiel, and E. Kasneci. 2015. ExCuSe: Robust Pupil Detection in Real-World Scenarios. In *16th International Conference on Computer Analysis of Images and Patterns (CAIP 2015)*.
- [47] Wolfgang Fuhl, Thomas C Kübler, Thiago Santini, and Enkelejda Kasneci. 2018. Automatic Generation of Saliency-based Areas of Interest for the Visualization and Analysis of Eye-tracking Data.. In *VMV*. 47–54.
- [48] Wolfgang Fuhl, Yao Rong, and Kasneci Enkelejda. 2020. Fully Convolutional Neural Networks for Raw Eye Tracking Data Segmentation, Generation, and Reconstruction. In *Proceedings of the International Conference on Pattern Recognition*. 0–0.
- [49] Wolfgang Fuhl, Yao Rong, Thomas Motz, Michael Scheidt, Andreas Hartel, Andreas Koch, and Enkelejda Kasneci. 2020. Explainable Online Validation of Machine Learning Models for Practical Applications. In *Proceedings of the International Conference on Pattern Recognition*. 0–0.
- [50] W. Fuhl, W. Rosenstiel, and E. Kasneci. 2019. 500,000 images closer to eyelid and pupil segmentation. In *Computer Analysis of Images and Patterns, CAIP*.
- [51] W Fuhl, N Sanamrad, and E Kasneci. 2021. The Gaze and Mouse Signal as additional Source for User Fingerprints in Browser Applications. *arXiv preprint arXiv:2101.03793* (01 2021).
- [52] W. Fuhl, T. Santini, D. Geisler, T. C. Kübler, and E. Kasneci. 2017. EyeLad: Remote Eye Tracking Image Labeling Tool. In *12th Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications (VISIGRAPP 2017)*.
- [53] W. Fuhl, T. Santini, D. Geisler, T. C. Kübler, W. Rosenstiel, and E. Kasneci. 2016. Eyes Wide Open? Eyelid Location and Eye Aperture Estimation for Pervasive Eye Tracking in Real-World Scenarios. In *ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct publication – PETMEI 2016*.

- [54] W. Fuhl, T. Santini, and E. Kasneci. 2017. Fast and Robust Eyelid Outline and Aperture Detection in Real-World Scenarios. In *IEEE Winter Conference on Applications of Computer Vision (WACV 2017)*.
- [55] Wolfgang Fuhl, Thiago Santini, and Enkelejda Kasneci. 2017. Fast camera focus estimation for gaze-based focus control. *arXiv preprint arXiv:1711.03306* (2017).
- [56] Wolfgang Fuhl, Thiago Santini, Gjergji Kasneci, and Enkelejda Kasneci. 2016. Pupilnet: Convolutional neural networks for robust pupil detection. *arXiv preprint arXiv:1601.04902* (2016).
- [57] Wolfgang Fuhl, Thiago Santini, Gjergji Kasneci, Wolfgang Rosenstiel, and Enkelejda Kasneci. 2017. Pupilnet v2. 0: Convolutional neural networks for cpu based real time robust pupil detection. *arXiv preprint arXiv:1711.00112* (2017).
- [58] W. Fuhl, T. Santini, T. C. Kübler, and E. Kasneci. 2016. ElSe: Ellipse Selection for Robust Pupil Detection in Real-World Environments. In *Proceedings of the Ninth Biennial ACM Symposium on Eye Tracking Research & Applications (ETRA)*. 123–130.
- [59] W. Fuhl, T. Santini, T. Kuebler, N. Castner, W. Rosenstiel, and E. Kasneci. 2018. Eye movement simulation and detector creation to reduce laborious parameter adjustments. *arXiv preprint arXiv:1804.00970* (2018).
- [60] W. Fuhl, T. Santini, C. Reichert, D. Claus, A. Herkommer, H. Bahmani, K. Rifai, S. Wahl, and E. Kasneci. 2016. Non-Intrusive Practitioner Pupil Detection for Unmodified Microscope Oculars. *Elsevier Computers in Biology and Medicine* 79 (12 2016), 36–44.
- [61] Wolfgang Fuhl, Marc Tonsen, Andreas Bulling, and Enkelejda Kasneci. 2016. Pupil detection for head-mounted eye tracking in the wild: An evaluation of the state of the art. In *Machine Vision and Applications*. 1–14.
- [62] Wolfgang Fuhl, Marc Tonsen, Andreas Bulling, and Enkelejda Kasneci. 2016. Pupil detection for head-mounted eye tracking in the wild: an evaluation of the state of the art. *Machine Vision and Applications* 27, 8 (2016), 1275–1288.
- [63] Stephan Joachim Garbin, Oleg Komogortsev, Robert Cavin, Gregory Hughes, Yiru Shen, Immo Schuetz, and Sachin S Talathi. 2020. Dataset for Eye Tracking on a Virtual Reality Platform. In *ACM Symposium on Eye Tracking Research and Applications*. 1–10.
- [64] Kaiping He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 770–778.
- [65] Christine Hooker and Sohee Park. 2005. You must be looking at me: The nature of gaze perception in schizophrenia patients. *Cognitive neuropsychiatry* 10, 5 (2005), 327–345.
- [66] Thomas E Hutchinson, K Preston White, Worthy N Martin, Kelly C Reichert, and Lisa A Frey. 1989. Human-computer interaction using eye-gaze input. *IEEE Transactions on systems, man, and cybernetics* 19, 6 (1989), 1527–1534.
- [67] Sergey Ioffe and Christian Szegedy. 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167* (2015).
- [68] Yukako Ishiyama, Hiroshi Murata, and Ryo Asaoka. 2015. The usefulness of gaze tracking as an index of visual field reliability in glaucoma patients. *Investigative ophthalmology & visual science* 56, 11 (2015), 6233–6236.
- [69] Enkelejda Kasneci, Katrin Sippel, Kathrin Aehling, Martin Heister, Wolfgang Rosenstiel, Ulrich Schiefer, and Elena Papageorgiou. 2014. Driving with binocular visual field loss? A study on a supervised on-road parcours with simultaneous eye and head tracking. *PloS one* 9, 2 (2014), e87470.
- [70] Joohwan Kim, Michael Stengel, Alexander Majercik, Shalini De Mello, David Dunn, Samuli Laine, Morgan McGuire, and David Luebke. 2019. Nvgaze: An anatomically-informed dataset for low-latency, near-eye gaze estimation. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–12.
- [71] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [72] Rakshit Kothari, Zhizhuo Yang, Christopher Kanan, Reynold Bailey, Jeff B Pelz, and Gabriel J Diaz. 2020. Gaze-in-wild: A dataset for studying eye and head coordination in everyday activities. *Scientific reports* 10, 1 (2020), 1–18.
- [73] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. 2015. Deep learning. *nature* 521, 7553 (2015), 436–444.
- [74] Matthew Luckiesh and Frank K Moss. 1937. The eyelid reflex as a criterion of ocular fatigue. *Journal of experimental Psychology* 20, 6 (1937), 589.
- [75] Christopher D McMurrough, Vangelis Metsis, Jonathan Rich, and Fillia Makedon. 2012. An eye tracking dataset for point of gaze detection. In *Proceedings of the Symposium on Eye Tracking Research and Applications*. 305–308.
- [76] Anjul Patney, Marco Salvi, Joohwan Kim, Anton Kaplanyan, Chris Wyman, Nir Benty, David Luebke, and Aaron Lefohn. 2016. Towards foveated rendering for gaze-tracked virtual reality. *ACM Transactions on Graphics (TOG)* 35, 6 (2016), 179.
- [77] P Jonathon Phillips, Kevin W Bowyer, and Patrick J Flynn. 2007. Comments on the CASIA version 1.0 iris data set. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 29, 10 (2007), 1869–1870.
- [78] Hugo Proença and Luís A Alexandre. 2005. UBIRIS: A noisy iris image database. In *International Conference on Image Analysis and Processing*. Springer, 970–977.
- [79] Hugo Proença, Silvio Filipe, Ricardo Santos, Joao Oliveira, and Luis A Alexandre. 2009. 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 (2009), 1529–1535.
- [80] Dale Purves, George J Augustine, David Fitzpatrick, Lawrence C Katz, Anthony-Samuel LaMantia, James O McNamara, S Mark Williams, et al. 2001. Types of eye movements and their functions. *Neuroscience* (2001), 361–390.
 - [81] Ning Qian. 1999. On the momentum term in gradient descent learning algorithms. *Neural networks* 12, 1 (1999), 145–151.
 - [82] Yashas Rai, Jesús Gutiérrez, and Patrick Le Callet. 2017. A dataset of head and eye movements for 360 degree images. In *Proceedings of the 8th ACM on Multimedia Systems Conference*. 205–210.
 - [83] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. 2015. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*. Springer, 234–241.
 - [84] Frank Rosenblatt. 1958. The perceptron: a probabilistic model for information storage and organization in the brain. *Psychological review* 65, 6 (1958), 386.
 - [85] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. 2018. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 4510–4520.
 - [86] Thiago Santini, Wolfgang Fuhl, Thomas Kübler, and Enkelejda Kasneci. 2016. Bayesian identification of fixations, saccades, and smooth pursuits. In *Proceedings of the Ninth Biennial ACM Symposium on Eye Tracking Research & Applications*. 163–170.
 - [87] Bernhard Schölkopf, Alexander J Smola, Francis Bach, et al. 2002. *Learning with kernels: support vector machines, regularization, optimization, and beyond*. MIT press.
 - [88] Virginia E Sturm, Megan E McCarthy, Ira Yun, Anita Madan, Joyce W Yuan, Sarah R Holley, Elizabeth A Ascher, Adam L Boxer, Bruce L Miller, and Robert W Levenson. 2011. Mutual gaze in Alzheimer’s disease, frontotemporal and semantic dementia couples. *Social Cognitive and Affective Neuroscience* 6, 3 (2011), 359–367.
 - [89] Lech Świrski, Andreas Bulling, and Neil Dodgson. 2012. Robust real-time pupil tracking in highly off-axis images. In *Proceedings of the Symposium on Eye Tracking Research and Applications*. 173–176.
 - [90] Lech Świrski and Neil Dodgson. 2013. A fully-automatic, temporal approach to single camera, glint-free 3d eye model fitting. *Proc. PETMEI* (2013), 1–11.
 - [91] Lech Świrski and Neil Dodgson. 2014. Rendering synthetic ground truth images for eye tracker evaluation. In *Proceedings of the Symposium on Eye Tracking Research and Applications*. 219–222.
 - [92] Tieniu Tan, Zhaofeng He, and Zhenan Sun. 2010. Efficient and robust segmentation of noisy iris images for non-cooperative iris recognition. *Image and vision computing* 28, 2 (2010), 223–230.
 - [93] Marc Tonsen, Xucong Zhang, Yusuke Sugano, and Andreas Bulling. 2016. Labelled pupils in the wild: a dataset for studying pupil detection in unconstrained environments. In *Proceedings of the Ninth Biennial ACM Symposium on Eye Tracking Research & Applications*. 139–142.
 - [94] Peter Wild, James Ferryman, and Andreas Uhl. 2015. Impact of (segmentation) quality on long vs. short-timespan assessments in iris recognition performance. *IET Biometrics* 4, 4 (2015), 227–235.
 - [95] Erroll Wood, Tadas Baltrusaitis, Xucong Zhang, Yusuke Sugano, Peter Robinson, and Andreas Bulling. 2015. Rendering of eyes for eye-shape registration and gaze estimation. In *Proceedings of the IEEE International Conference on Computer Vision*. 3756–3764.
 - [96] Zhengyang Wu, Srivignesh Rajendran, Tarrence van As, Joelle Zimmermann, Vijay Badrinarayanan, and Andrew Rabinovich. 2020. MagicEyes: A Large Scale Eye Gaze Estimation Dataset for Mixed Reality. *arXiv preprint arXiv:2003.08806* (2020).