What Do You See in Vehicle? Comprehensive Vision Solution for In-Vehicle Gaze Estimation

Yihua Cheng[1]   Yaning Zhu[2]   Zongji Wang[3]   Hongquan Hao[4]   Yongwei Liu[4]

Shiqing Cheng[4]   Xi Wang[4]   Hyung Jin Chang[1]

University of Birmingham[1],   Huazhong University of Science and Technology[2]

NIST, Chinese Academy of Sciences[3],   ClamCar[4]

y.cheng.2@bham.ac.uk, h.j.chang@bham.ac.uk

Abstract

Driver's eye gaze holds a wealth of cognitive and intentional cues crucial for intelligent vehicles. Despite its significance, research on in-vehicle gaze estimation remains limited due to the scarcity of comprehensive and well-annotated datasets in real driving scenarios. In this paper, we present three novel elements to advance in-vehicle gaze research. Firstly, we introduce IVGaze, a pioneering dataset capturing in-vehicle gaze, collected from 125 subjects and covering a large range of gaze and head poses within vehicles. In this dataset, we propose a new vision-based solution for in-vehicle gaze collection, introducing a refined gaze target calibration method to tackle annotation challenges. Second, our research focuses on in-vehicle gaze estimation leveraging the IVGaze. In-vehicle face images often suffer from low resolution, prompting our introduction of a gaze pyramid transformer that leverages transformer-based multilevel features integration. Expanding upon this, we introduce the dual-stream gaze pyramid transformer (GazeDPTR). Employing perspective transformation, we rotate virtual cameras to normalize images, utilizing camera pose to merge normalized and original images for accurate gaze estimation. GazeDPTR shows state-of-the-art performance on the IVGaze dataset. Thirdly, we explore a novel strategy for gaze zone classification by extending the GazeDPTR. A foundational tri-plane and project gaze onto these planes are newly defined. Leveraging both positional features from the projection points and visual attributes from images, we achieve superior performance compared to relying solely on visual features, substantiating the advantage of gaze estimation. Our project is available at https://yihua.zone/work/ivgaze.

1. Introduction

Understanding driver intention and behavior based on driver gaze is in high demand in intelligent vehicles, facilitating
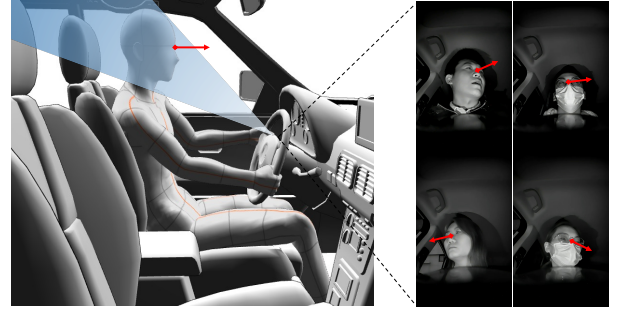

Figure 1. In-vehicle gaze estimation illustration. The driver's gaze direction is estimated based on the facial images captured by the camera behind the steering wheels.

diverse applications such as in-vehicle interaction [1, 2, 25] and driver monitor systems [17, 21, 22]. Recent advances in vehicle gaze estimation concentrate primarily on gaze zone estimation [16, 18, 31, 32]. These approaches define multiple coarse regions, such as side mirrors and windshields, and conduct classification based on face images.

Gaze estimation[1] serves as an upstream task of gaze zone estimation and can offer more precise information to understand driver attention. However, these methods typically require a large-scale dataset for training. Although there are numerous gaze datasets, collected in indoor [39, 40] or outdoor [20] environments, their applicability to the vehicle environment is limited due to the different environments and camera settings, resulting in suboptimal performance. Creating an in-vehicle gaze dataset proves challenging due to the confined and irregular nature of the vehicular environment. Constructing in-vehicle gaze collection systems remains an unsolved issue, as traditional gaze collection systems are impractical for use within vehicles. The absence of in-vehicle gaze datasets acts as a significant barrier to the progress of in-vehicle gaze estimation.

In this paper, we present a comprehensive vision-based

---

[1]Our work focuses on gaze direction estimation. We abbreviate gaze direction as gaze in the rest.

in-vehicle gaze estimation research: a novel vision-based gaze collection system for vehicles, offering a first-of-its-kind in-vehicle gaze dataset, a dual-stream gaze pyramid transformer for accurate in-vehicle gaze estimation, and its extension to gaze zone classification, showcasing its effectiveness in enhancing gaze estimation.

First, we introduce IVGaze, an in-vehicle gaze dataset collected from 125 subjects. IVGaze provides a dense distribution of gaze directions, covering a wide range within in-vehicle environments. It contains various conditions, including diverse head poses, eye movements, illumination variations, and the presence of face accessories such as glasses, sunglasses, and masks. To collect IVGaze, we propose a vision-based gaze collection system. The system does not require dedicated eye-tracking devices and is easy to reproduce. We use a single camera for facial appearance capture and paste stickers in vehicles as gaze targets. However, a significant challenge lies in calibrating the 3D position of gaze targets. This challenge arises because gaze targets are out of the camera's field of view (out-of-FoV). To address this issue, we present a refined gaze target calibration method, which leverages an auxiliary camera to capture gaze targets and a transparent chessboard for calibration.

Secondly, we explore in-vehicle gaze estimation using the IVGaze. Face images often suffer from low resolution due to the inherent limitations of cameras in vehicles. We propose a gaze pyramid transformer that utilizes a transformer to integrate multilevel features. Expanding upon this, we propose a dual-stream gaze pyramid transformer (GazeDPTR). We rotate virtual cameras via perspective transformation to normalize images, and leverage camera pose to merge normalized and original images. GazeDPTR shows state-of-the-art results on IVGaze.

Thirdly, we extend GazeDPTR for the downstream gaze zone classification task. It is challenging to compute the intersection of gaze and the vehicle. We define a foundational tri-plane and project gaze to the tri-plane. We extract positional features from intersection points and predict gaze zones using both positional features and visual features from face images. Our experiment demonstrates that the gaze zone classification can be further enhanced by positional features, showing the advantage of gaze estimation.

2. Related Works

2.1. Gaze Data Collection

The human gaze is inherently implicit and poses a challenge for objective measurement, making gaze annotation complex in gaze data collection. Some methods capture the human gaze through intrusive devices such as eye-tracking glasses [13, 19]. However, the eye-tracking glasses have a notable impact on the quality of the captured facial images. Fischer et al. [13] attempted to mitigate this impact

Table 1. The comparison of gaze estimation datasets. Gaze annotation is challenging in the vehicle environment, which results in existing in-vehicle datasets only providing gaze zone annotation. Our work addresses this issue and contributes the first in-vehicle gaze dataset containing gaze annotation and natural face images.

Datasets	Env.	# Sub.	Annotation	
			Gaze	Zone
Gaze360 [20]	Outdoor	238	✓	-
MPIIGaze [39]		15	✓	-
EyeDiap [15]	Indoor	16	✓	-
EVE [27]		54	✓	-
ETH-XGaze [40]		110	✓	-
Vora et al. [31]		10	✗	✓
Jha & Busso [18]		16	✗	✓
Wang et al. [32]	Vehicle	3	✗	✓
Rangesh et al. [28]		13	✗	✓
Ghosh et al. [16]		338	✗	✓
IVGaze (Ours)		125	✓	✓

using GAN, but there are also some artifacts in the resulting face images. Vision-based gaze collection systems typically define gaze as direction vectors originating from facial centers towards specific gaze targets [27, 39, 40]. However, these gaze targets often remain outside the camera's field of view. This out-of-field view challenge complicates the calibration of gaze target positions. Zhang et al. [39, 40] sets screen points as gaze targets and uses a mirror to calibrate the screen plane. Kellnhofer et al. [20] uses 360° panoramic cameras to capture gaze targets and human faces simultaneously where the panoramic camera is pre-calibrated. However, these strategies are not applicable to vehicles.

In-vehicle gaze dataset usually defines different region such as windshield and left/right mirror in vehicles and perform gaze zone classification [12, 14, 16, 18, 23, 30, 34]. Kasahara et al. [19] collects an in-vehicle gaze dataset but subjects are required to wear eye-tracking glasses, which means the dataset is not applicable in the real world. Our gaze collection system does not require dedicated devices and produces natural face images.

2.2. Appearance-based Gaze Estimation

Appearance-based gaze estimation directly learns mapping function from facial appearance to human gaze [10]. Conventional gaze estimation methods extract eye features from eye images [7, 9, 26, 36]. They concatenate the head pose vector with eye features for gaze estimation. Recent methods directly learn gaze from face images [3, 4, 11, 33, 35]. Face images provide both eye region and head pose information, resulting in superior performance compared to methods relying solely on eye images. However, the subtlety of the human eye in face images poses a challenge as it can be easily overlooked by the network. Zhang et al. [37] utilize a learnable attention map to guide the net-
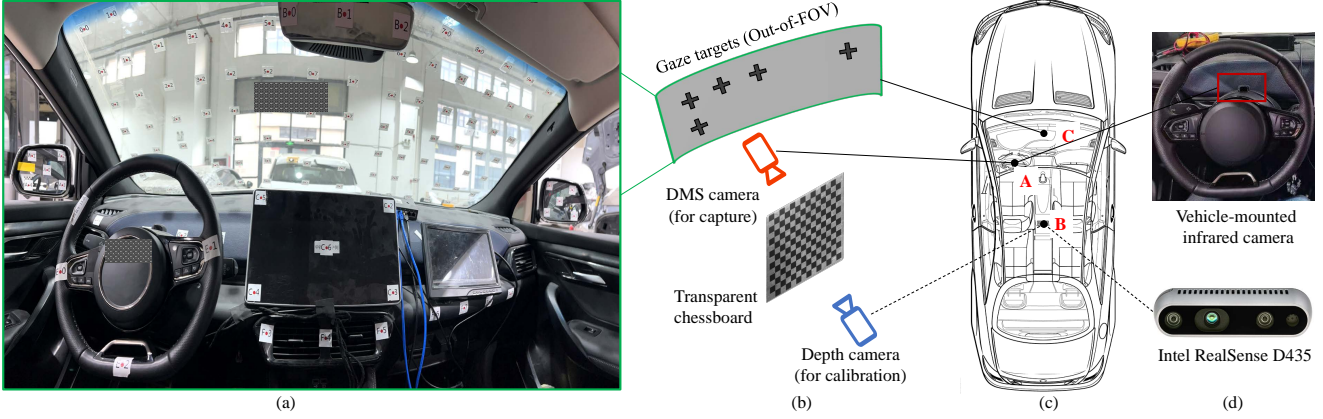
Figure 2. We construct a vision-based in-vehicle gaze collection system comprising a DMS camera, a depth camera, and strategically placed gaze targets, as depicted in (c). The DMS camera is positioned behind the steering wheel to capture drivers' facial appearances, while the gaze targets, positioned beyond the DMS camera's field-of-view (FoV), such as on the windshield, remain unobserved. The depth camera, utilized for calibration purposes, is temporarily installed for capturing gaze target positions in 3D with respect to its own coordinates, and it is removed during data collection. To facilitate the calibration of the depth camera's pose relative to the DMS camera, we propose employing a *transparent chessboard*, which is placed between the two cameras.

work's focus specifically on the eye. Cheng *et al*. [11] crop eye images from face images and construct a cascade network to leverage them. In recent developments, transformers have showcased remarkable capabilities in gaze estimation. Cheng *et al*. [5] demonstrate that transformers can achieve state-of-the-art performance across various benchmarks. They apply transformers to address dual-view gaze estimation and show that dual-view images outperform single-view images [6].

3. In-Vehicle Gaze Data Collection System

In this section, we introduce our vision-based system for in-vehicle gaze data collection, eliminating the need for dedicated eye-tracking devices. Our system includes a neat gaze target calibration method that effectively addresses the core challenge in gaze annotation. Leveraging this system, we collect the first in-vehicle gaze dataset from 125 subjects. The dataset spans a diverse range of gaze, facilitating and advancing future research in in-vehicle gaze estimation.

3.1. System Setup

We place gaze targets within the vehicle, and subjects are instructed to look at each designated target during the data collection process. Simultaneously, we record their facial appearance along with corresponding gaze targets.

Camera. Our system uses one camera of a driver monitor system (DMS) for facial appearance capture. The DMS camera is an infrared camera with a capture resolution of 1280×800. The camera is located behind the steering wheel and directly points at the head region of drivers.

Gaze targets. We mark gaze targets with red points on stickers, and each target is uniquely labeled with a distinct number printed on stickers to aid differentiation. We strate-

gically position these stickers within various gaze regions within the vehicle, including the windshield, left and right-side mirrors, rear-view mirror, center console, speedometer, handbrake, and etc. These targets cover a large gaze region, satisfying the gaze estimation requirements in vehicles.

3.2. Collection Procedure

We meticulously design the collection procedure to minimize error. During the collection process, participants are instructed to look at specific gaze targets in a predefined sequence. They are also required to speak the corresponding number associated with each target to confirm their focus accuracy. Participants must maintain focus on the target for 3 seconds to facilitate image capture. We conduct a preliminary check on the captured images and discard any erroneous ones. Once this verification is complete, the current data collection phase concludes, and participants are instructed to shift their gaze to the next gaze targets.

We design three postures for each participant. Participants should complete the collection process three times with different postures. The first posture requires participants not to change head pose. They should preserve the same head pose and only move eyeballs to focus a gaze target. We aim to collect sufficient eye movement data with this posture. We do not require participants to focus on the target where participants think it is hard to focus with only eye movement, *e.g*., the target in the side mirror. Participants can rotate their heads to look at these points but they are also required to preserve the new head pose until they cannot look at any targets with the current head pose. The second posture requires participants not to move their head position but they can rotate their heads to focus a gaze target. We do not set any constraints in the third posture. Participants can freely perform head movements in this pos-

(a) IVGaze dataset samples   (b) Dataset statistics across illumination intensity   (c) Dataset statistics across face accessories
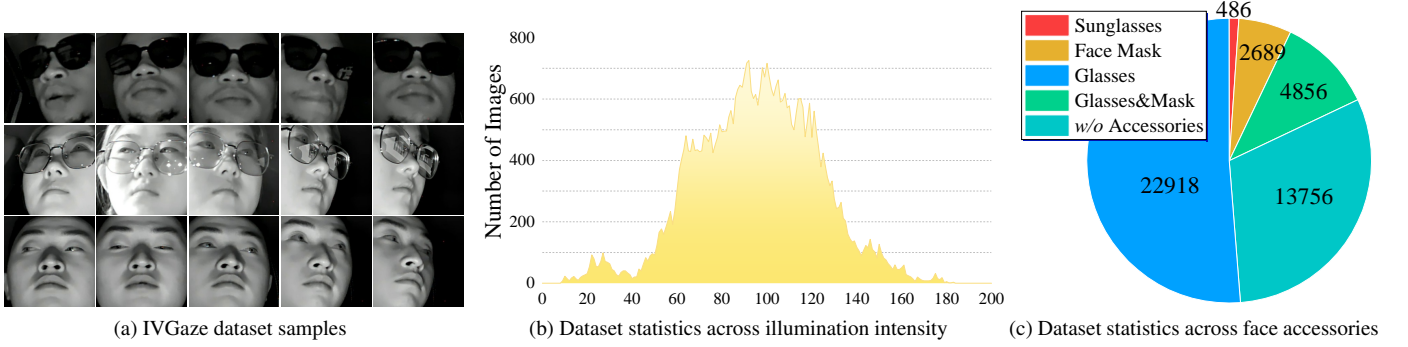
Figure 3. Our dataset is collected using IR cameras in the vehicle environment. (a) We present image samples of IVGaze, highlighting the challenges posed by realistic in-vehicle conditions, including cases with sunglasses and reflections in glasses. (b) We categorize the image count based on their mean pixel value, showing the diversity of illumination conditions. (c) The image count is analyzed based on face accessories including glasses, sunglasses, and masks.

ture. The last two postures bring head pose variation for IVGaze. We do not use devices such as headrests to constrain participants. Participants only need to preserve the specific posture and focus on targets in a comfortable way.

3.3. Gaze Annotation

Vision-based gaze collection systems define gaze as unit direction vectors originating from the face center $\mathbf{o}$ and extending toward gaze targets $\mathbf{t}$ [10]. Our system follows this framework and the gaze annotation can be computed as $\mathbf{g} = (\mathbf{t} - \mathbf{o})/||\mathbf{t} - \mathbf{o}||_2$. The gaze annotation is decomposed into 3D face center estimation and gaze target calibration.

3D face center estimation has been effectively addressed in previous research. We first detect facial landmarks and fit a 3D morphable face model for 3D facial landmarks [24, 40]. We select the 3D position of the nose as gaze origin $\mathbf{o}$. However, the calibration of gaze target positions presents a challenge because gaze targets are out-of-FoV. It means that the DMS camera cannot capture any images of the gaze targets, complicating the calibration problem.

3.4. Gaze Target Calibration

We need the 3D gaze target position $\mathbf{t}$ for gaze annotation. However, a key problem is all gaze targets are out-of-FoV. Conventional gaze datasets often present gaze targets on a screen and employ a mirror to solve the out-of-FoV problem, where the mirror reflects the content of the screen [36]. However, these methods are not suitable for vehicle scenarios and can lead to a substantial accumulation of errors.

In this section, we propose a neat and efficient gaze target calibration method. Our basic idea is to employ an auxiliary camera to capture gaze targets. This allows a straightforward calculation of the 3D gaze target positions w.r.t. the auxiliary camera coordinate system. However, it also introduces a new challenge: *how to calibrate extrinsic matrix of the auxiliary cameras w.r.t. DMS camera*? Stereo calibration is typically used to compute the pose of two cam-

eras. However, it requires two captured images sharing corresponding points. In our system, gaze targets are located behind the DMS camera which means the DMS camera and auxiliary camera should be oriented in opposite directions, making traditional stereo calibration infeasible.

To solve this problem, we propose to use a *transparent chessboard* for the pose calibration. We set a transparent chessboard between the two cameras and the two cameras respectively capture each side of the chessboard. We first calibrate the two camera poses w.r.t. the chessboard coordinate system and then compute the pose between two cameras. It is worth noting that there are two different chessboard coordinate systems corresponding to each side of the chessboard. We can derive the transformation matrix between the two chessboard coordinate systems where $\mathbf{R}_{\text{chess}} = \text{diag}\,(0, 0, -1)$ and $\mathbf{t}_{\text{chess}} = (0, 0, -d)$. The d is the thickness of the chessboard.

In detail, we use an Intel RealSense D435 depth camera as the auxiliary camera since it can provide accurate 3D positions. We show the camera layout in Fig. 2(b). We can obtain the transformation matrix $\mathbf{R}_{\text{dms}}, \mathbf{t}_{\text{dms}}$ and $\mathbf{R}_{\text{depth}}, \mathbf{t}_{\text{depth}}$ which can transfer a point from chessboard coordinate systems to camera coordinate systems. We have

$$\mathbf{R}_{\text{rot}} = \mathbf{R}_{\text{dms}}\mathbf{R}_{\text{chess}}\mathbf{R}_{\text{depth}}^{-1}, \tag{1}$$

$$\mathbf{t}_{\text{rot}} = -\mathbf{R}_{\text{dms}}\mathbf{R}_{\text{chess}}\mathbf{R}_{\text{depth}}^{-1}\mathbf{t}_{\text{depth}} + \mathbf{R}_{\text{dms}}\mathbf{t}_{\text{chess}} + \mathbf{t}_{\text{dms}}. \tag{2}$$

We use these matrices to convert points from the depth camera coordinate system into the DMS camera coordinate system. Please refer to the supplementary material for details.

3.5. In-Vehicle Gaze Dataset

We collect the first in-vehicle gaze dataset IVGaze which provides dense gaze annotation and natural face images . We show the dataset samples in Fig. 3 (a).

Dataset Statistics. We collect 44,705 images from 125 subjects. The number of images per subject ranges from
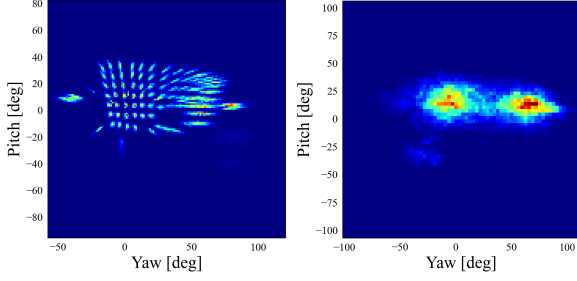
Figure 4. We show the distribution of data for gaze (left) and head movements (right). Brighter regions denote higher data density.

169 to 946. Most subjects provide around 300 images per person, *i.e.*, they repeat the collection procedure three times with three different postures. 40 subjects are female and 85 subjects are male, with ages spanning from 20 to 50 years.

Collection Conditions. The dataset was collected between 9 am and 7 pm in outdoor environments, covering a wide range of lighting conditions. We also collect images in indoor environments, *i.e.*, a garage. The dataset statistics across pixel intensity are shown in Fig. 3 (b).

Face Accessories. We consider two face accessories during the collection: glasses and masks. For subjects who did not wear glasses, we provided glasses and asked them to repeat the collection wearing glasses. We provided white and black masks for some subjects and asked them to repeat the collection wearing masks. We also required a few subjects to wear sunglasses to facilitate future research. We show the statistics in Fig. 3 (c).

Data Distribution. We plot the angular distribution of gaze and head pose. Note that the data normalization method (introduced in the next section) changes gaze directions and head poses with a rotation matrix. Fig. 4 shows the distribution in normalization space. The horizontal gaze is from $-50°$ to $90°$, and the vertical gaze is from $-40°$ to $40°$. It is worth noting that the vehicle environment has limited space, and our dataset already covers almost all regions. Our dataset is collected in the real vehicle environment, and the driver seat in the vehicle is located on the left. Therefore, our dataset has a relatively small distribution on the left. The right figure in Fig. 4 shows the head pose distribution. Our dataset contains a large range of head poses in the yaw axis, as subjects need to look at gaze targets from the left-side mirror to the right-side mirror.

4. In-Vehicle Gaze Estimation

The in-vehicle environment brings new settings and challenges. In this section, we systematically explore in-vehicle gaze estimation. We introduce a revised data preprocessing method and describe a novel gaze estimation network. We also extend the network for an application and prove the advantage of in-vehicle gaze estimation.

4.1. Revisiting Data Normalization in Vehicle

Data normalization rotates and translates cameras via perspective transformation in face images to reduce head pose variation. Conventional methods usually define the rotated camera coordinate system $\mathbf{C}_r$ based on head pose [29, 38]. However, these methods cannot bring performance improvement in the vehicle environment.

The x-axis of $\mathbf{C}_r$ is typically defined as the x-axis of head coordinate systems. It means they will rotate images to keep the head straight. We observed that the designed x-axis will produce unstable results, especially in the extreme head pose. Therefore, we propose that do not rotate virtual cameras based on the x-axis of head coordinate systems. In detail, we compute the z-axis of $\mathbf{C}_r$ as direction vectors from cameras to face centers. We use the x-axis of the virtual camera, *i.e.*, $(1, 0, 0)$ rather than the head coordinate system. We compute the y-axis as $y = z \times x$ and also recompute the x-axis as $x = y \times z$. Finally, we have a rotation matrix $\mathbf{R} = [x; y; z]$ and use a scale matrix S to maintain the user-camera distance. We warp the image based on matrix $\mathbf{SR}$. The human gaze is also changed as $\mathbf{g}^n = \mathbf{R}\mathbf{g}^o$, where $\mathbf{g}^n$ and $\mathbf{g}^o$ denotes gaze in the normalization and origin space.

4.2. Gaze Pyramid Transformer

Gaze estimation methods typically learn gaze from face images. However, face images captured in vehicles often suffer from low resolution owing to the limited capabilities of the camera. This low resolution can lead conventional methods to overlook subtle details, especially in the eye region. Some techniques address this challenge by cropping eye images and processing them separately [8], but these approaches encounter difficulties when faced with extreme head poses, as one of the eye images may not be visible.

We propose a gaze pyramid transformer (GazePTR) in this section. We build a feature pyramid and input multi-level features into a transformer for gaze estimation. As shown in Fig. 5, our network uses a convolution network to extract features from facial images. We collect feature maps from different levels rather than the last level. We use 1×1 convolution layers and global average pooling layers to preserve the same feature dimension. Finally, we input the multi-level feature into a transformer. A learnable token is used to aggregate multi-level features for gaze estimation.

However, the network has the same performance achieved by directly utilizing the last feature map for gaze estimation. This result can be attributed to the network's tendency to overlook low-level features. To address this issue, we introduce additional constraints by estimating gaze from each level of features and calculating corresponding loss. It can be understood that we first extract multi-level features, with each feature being potentially useful. A transformer is used to effectively ensemble these features.
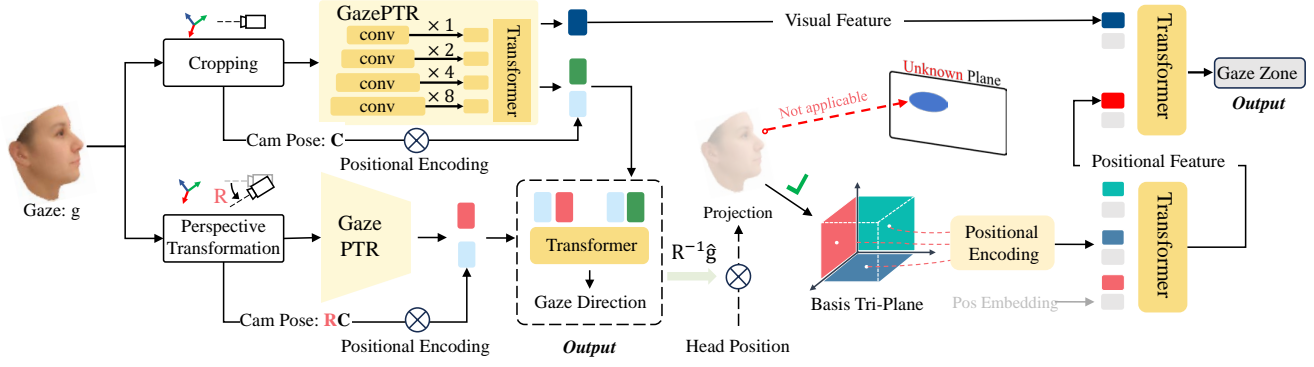
Figure 5. The GazeDPTR directly crop face for origin images and rotates virtual cameras via perspective transformation for normalized images. It builds a dual-stream network to extract features from the two images based on the GazePTR for feature extraction which integrates multi-level features via transformers. To further merge the features from two streams, we leverage a transformer where camera pose is used as the positional feature in the transformer. We define the original camera pose as $\mathbf{C} = diag(1,1,1)$ and the camera pose in normalization space is $\mathbf{RC}$. We also extend the network for gaze zone classification. We define a tri-plane and project gaze into them. We extract positional features from three intersection points via a transformer. We also extract visual features from images and predict the gaze zone based on both visual features and positional features. The whole network is trained in an end-to-end manner.

4.3. Dual-Stream Gaze Pyramid Transformer

Dual-view images have been demonstrated to outperform single-view images in gaze estimation [6]. In our work, we rotate cameras in data normalization through perspective transformation. This inspired us to formulate the dual-camera setting leveraging both original and normalized images. However, dual-view images improve accuracy since they contain more visual information. It remains uncertain whether the combination of normalized and original images can provide additional insights beyond what each offers.

To validate our hypothesis, we executed an oracle baseline. We separately train GazePTR on the original and the normalized datasets, selecting the best result from each image pair. The result reveals a significant performance improvement, validating the inherent advantages of utilizing both images. Please refer to the supplementary for details.

Therefore, we propose a dual-stream gaze pyramid transformer (GazeDPTR). Our method leverages both normalized images and original images for gaze estimation. We employ GazePTR to extract features from the two images and a transformer is used to integrate the two features for gaze estimation. Note that, the two features are in different camera coordinate systems and correspond to different gaze. To establish the connection, we use the camera pose as the positional information in the transformer. However, it is impossible the calibrate a virtual camera. We define the camera pose of original images as $\mathbf{C} = diag(1,1,1)$ and the camera pose of normalized images is $\mathbf{RC}$. We extract the z-axis and use positional encoding for camera pose [6].

4.4. Strategy for Gaze Zone Classification

Detecting driver attention is essential in driver monitoring systems. In this section, we introduce a novel strategy to extend GazeDPTR for gaze zone classification. GazeDPTR estimates gaze direction from face images. While the conventional solution involves projecting gaze direction and computing intersections within the vehicle, it is not practical without a 3D vehicle model. Zhang *et al*. use a tri-plane for self-supervised loss function [35]. Our alternative strategy defines a foundational tri-plane and computes the intersection of gaze vectors with this tri-plane. We extract positional features from three intersection points to facilitate gaze zone classification.

In detail, we define a tri-plane with normal vectors $(1,0,0)$, $(0,1,0)$, and $(0,0,1)$, intersecting at the origin $(0,0,0)$. Given the estimated gaze vector $\mathbf{g}^n$, we convert it back to the original space using $\mathbf{g}^o = \mathbf{R}^{-1}\mathbf{g}^n$. We then project $\mathbf{g}^o$ onto the tri-plane with the face center $\mathbf{o}$ and obtain three intersection points. To enhance the positional features, we apply positional encoding to these points and input them into a 2-layer transformer. The positional encoding not only increases feature dimensions but also normalizes them within the range of -1 to 1. Additionally, we use another learnable token to aggregate visual features from the original images. Gaze zone classification is performed using both visual features and positional features.

We train the whole network in an end-to-end manner. We use L1 loss for gaze estimation and cross-entropy loss for gaze zone classification. Please refer to the supplementary for implementation details.

5. Experiment

Dataset: IVGaze contains 44,705 images of 125 subjects. To perform within-dataset evaluation, we divide the dataset into three subsets based on subjects. The image numbers of the three subsets are 15,165, 14,674, and 14,866. We

Table 2. Performance comparison in gaze estimation. Our methods achieve better performance than the compared methods.

	Angular Error	Average Precision (AP)			
		$<2°$	$<4°$	$<6°$	$<8°$
FullFace [37]	$13.67°$	2.3%	8.8%	17.8%	28%
DWG [16]	$8.82°$	6.6%	21.7%	38.1%	53.2%
Gaze360 [20]	$8.15°$	9.2%	27.3%	44.6%	58.9%
FullFace$^+$ [37]	$7.48°$	14.2%	31.1%	46.7%	63.1%
GazeTR [5]	$7.33°$	17%	32.8%	47.5%	64.7%
XGaze [40]	$7.06°$	11.7%	32.7%	51.5%	66.7%
GazePTR	$7.04°$	17.6%	34%	49.3%	66.7%
GazeDPTR	$6.71°$	22.1%	36%	50.3%	68.4%

Table 3. We show the impacts of different face accessories on performance. The sunglasses bring a significant performance drop.

	Glasses		Mask		Sunglasses
	with	*w/o*	with	*w/o*	with
FullFace [37]	$14.43°$	$12.40°$	$15.20°$	$13.35°$	$21.39°$
DWG [16]	$9.20°$	$8.19°$	$9.43°$	$8.69°$	$17.43°$
Gaze360[20]	$8.30°$	$7.91°$	$8.95°$	$7.99°$	$17.99°$
FullFace$^+$ [37]	$7.59°$	$7.30°$	$8.37°$	$7.30°$	$16.50°$
XGaze [40]	$7.07°$	$7.03°$	$7.80°$	$6.90°$	$15.15°$
GazeTR [5]	$7.40°$	$7.22°$	$8.12°$	$7.17°$	$17.49°$
GazePTR	$7.13°$	$6.90°$	$7.78°$	$6.89°$	$16.54°$
GazeDPTR	$6.77°$	$6.63°$	$7.44°$	$6.57°$	$16.41°$

perform three-fold cross-validation on our dataset.

Evaluation Metric: We use the angular error as the evaluation metric of gaze estimation as most of the methods [10], where a smaller value represents a better model. However, we notice the angular error only shows the overall performance while cannot give deep insights. Therefore, we define the average precision (AP) where AP of $<k°$ means an estimation is considered correct if the angular error is lower than $k°$. Regarding the gaze zone classification, average precision is used as a common multi-class classification.

5.1. Comparison with SOTA Methods

We first compared our methods with SOTA methods in the in-vehicle gaze estimation task. We compared our methods with FullFace [37], Gaze360 [20], XGaze [40], DWG [16] and GazeTR [5]. We replaced the backbone of the FullFace method from AlexNet to ResNet18 for a more convincing comparison. We denote the new method as FullFace$^+$.

The result is shown in Tab. 2. GazePTR and GazeDPTR both show better performance than the compared methods. GazePTR and GazeTR have the same backbone and transformer architecture. However, GazePTR outperforms $0.29°$ thanGazeTR since GazePTR uses multi-level feature maps. GazeDPTR builds a dual-stream network and further brings $0.33°$ improvement than GazePTR. These results show the advantage of our methods.

Interestingly, XGaze has a similar angular error to GazePTR. However, it is significantly worse than GazePTR in AP of $<2°$. It is because that XGaze uses ResNet50 as the backbone. The deep backbone enhances the feature extraction ability but also easily overlooks the small eye region, making it challenging to achieve precise gaze estimation.

5.2. The Impact of Face Accessories

IVGaze provides rich samples of wearing face accessories, including glasses, masks, and sunglasses. We also evaluate their impacts on the accuracy of gaze estimation and show in Tab. 3. The result shows that eyeglasses have a
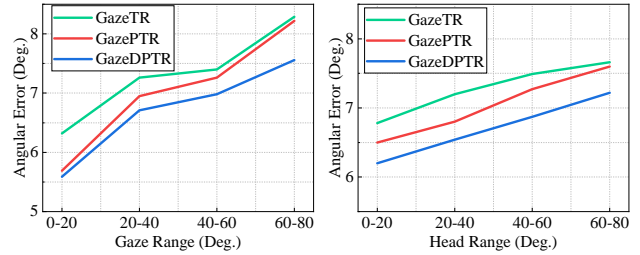

Figure 6. We compute the angular degree between gaze direction/head orientation with the frontal direction, *i.e.*, (0, 0, -1), and count the mean gaze estimation accuracy in different ranges. Interestingly, GazePTR brings significant performance improvement to GazeTR in the $0° - 20°$ gaze range. This is because subjects prefer to move their eyeball rather than their head when the gaze target is located at a close range. GazePTR utilizes multilevel features and can accurately capture eye movement.

large impact if the basic performance is not good. However, as the basic performance improves, the performance difference reduces. This is because most of glasses are transparent, enabling a robust model to extract features effectively. Sunglasses have a significant impact on performance, with average performance exceeding $15°$. Masks have a substantial impact on performance, as they obscure facial regions, complicating the extraction of head-related features.

5.3. Performance Distribution

We show the performance distribution to provide deep insight. We first show the performance distribution in different gaze ranges. We introduce a novel concept, the gaze range of $a°$-$b°$, which is defined as a set of samples satisfying the requirement that the angular degree between the gaze direction and the frontal gaze direction is in the range. The result is shown in Fig. 6. Interestingly, GazePTR and GazeDPTR have better performance in $0°$-$20°$. This is because subjects prefer to move their eye rather than head to look at objects in this range. Our methods leverage multilevel features and easily capture eye movement. We also show the performance distribution in different head ranges.

Table 4. We present the performance of GazePTR at each level of the feature. The higher-level feature achieves better performance. Our method integrates features from all levels and achieves the best performance.

	$1st$	$2nd$	$3rd$	$4th$	Ours
GazePTR	$10.33°$	$8.59°$	$7.61°$	$7.35°$	$7.04°$

Table 5. We present the performance of GazePTR separately training on the original images and normalized images. GazeDPTR leverages camera pose to establish a connection between the two images and use both two images for gaze estimation. We also removed the camera pose to perform an ablation study.

	Original images	Normalized images	GazeDPTR (*w/o* camera pose)	GazeDPTR
Acc	$7.44°$	$7.04°$	$7.03°$	$6.71°$

Our methods are robust in different head ranges where the maximum performance difference is only $1°$.

5.4. Ablation Study

GazePTR leverages multi-level features for gaze estimation. We first conduct an ablation study to demonstrate the advantage. We obtain the gaze performance of GazePTR at each level of feature and show it in Tab. 4. The result demonstrates a higher-level feature achieves better performance. GazePTR integrates multi-level features and achieves the best performance. The result proves the advantage of our design.

GazeDPTR leverages both original and normalized images along with their corresponding camera poses. To demonstrate the effectiveness of our design, we perform an ablation study in Tab. 5. Training GazePTR separately on original and normalized images yields performances of $7.44°$ and $7.04°$ respectively. GazeDPTR, which incorporates both image types, achieves a performance improvement of $0.34°$. We also experimented by excluding the positional encoding of the camera pose in GazeDPTR. Interestingly, the result is comparable to using only normalized images.

5.5. Additional Experiments

Data Normalization. We compare our method with the data normalization method [38] in the left of Tab. 6. The previous method cannot work well in vehicle environments while our method brings $0.3°$ performance improvement. **Case Study.** We visual gt and our prediction in Fig. 7.

5.6. Gaze Zone Classification

We propose a basis tri-plane to acquire positional feature and combine both positional feature and visual features for gaze zone classification. We respectively evaluate each feature and show the result in the right of Tab. 6. The visual feature achieves better performance than the positional feature since the positional feature is obtained from gaze pro-

Table 6. We evaluate different data normalization methods in the left table. Our method outperforms the previous method [38]. In the right table, we evaluate different features in the gaze zone classification. Our method uses both positional feature and visual features, achieving 2.4% improvement over visual features.

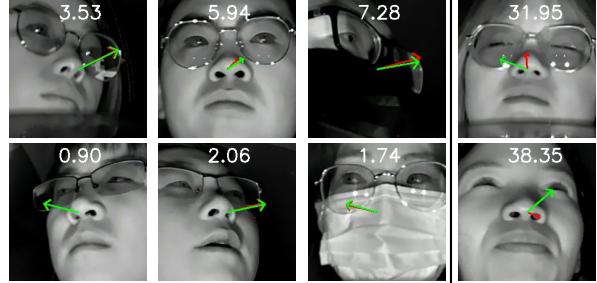| Normalization [38]	ours	GazeTR	GazePTR	Visual Feature	Positional Feature	AP
×	×	$7.77°$	$7.44°$	✓		79.4%
✓		$8.64°$	$8.53°$		✓	75.3%
	✓	$7.33°$	$7.04°$	✓	✓	81.8%

Figure 7. We illustrate GT with green lines and Prediction with red lines. The number shows the angular error in degree. The right two figures are failure cases since the eye is hard to capture.

jection. The projection increases the model interpretability but decreases the fitting ability. Our method uses the two features and gets 2.4% performance improvement. The result demonstrates that the gaze zone classification task could be enhanced by gaze direction cues, highlighting the advantage of our gaze estimation work.

6. Conclusion

In this work, we provide systematical research in in-vehicle gaze estimation, including dataset, algorithm, and extensive application. We first solve the data collection issue in a vehicle where a gaze target calibration method is proposed. We collect the first in-vehicle gaze dataset containing dense gaze annotation and natural face images. We further explore the algorithm for in-vehicle gaze estimation. Our work brings two deep insights 1) multi-level feature is useful to capture eye region information. 2) simultaneously leveraging original images and normalized images could achieve better performance. We also extend our work for the downstream gaze zone classification task. We demonstrate that the gaze direction cues could bring performance improvement for gaze zone classification.

7. Acknowledgments:

What Do You See in Vehicle? Comprehensive Vision Solution for In-Vehicle Gaze Estimation

Supplementary Material

Due to the page limitation, we present some details in the supplementary material. We first describe the details of method and then demonstrate more experiment results.

8. Methodology

8.1. Gaze Target Calibration

We uses a transparent chessboard for gaze target calibration. The mathematical deduction is shown in this section.

We set a transparent chessboard between the DMS camera and the depth camera. The two cameras both capture one side of the chessboard. Therefore, we can compute the pose matrices of the two cameras w.r.t. the chessboard coordinate system. We denote the pose matrices of the two cameras as $\{\mathbf{R}_{\mathrm{dms}}, \mathbf{t}_{\mathrm{dms}}\}$ and $\{\mathbf{R}_{\mathrm{depth}}, \mathbf{t}_{\mathrm{depth}}\}$. Given a point $\mathbf{p}_1$ in the chessboard coordinate system, we have

$$\mathbf{p}_{\mathrm{dms}} = \mathbf{R}_{\mathrm{dms}}\mathbf{p}_1 + \mathbf{t}_{\mathrm{dms}}. \tag{3}$$

Similarly, we can compute the 3D position $\mathbf{p}_{\mathrm{depth}}$ in the depth camera coordinate system with a given point $\mathbf{p}_2$.

$$\mathbf{p}_{\mathrm{depth}} = \mathbf{R}_{\mathrm{depth}}\mathbf{p}_2 + \mathbf{t}_{\mathrm{depth}}. \tag{4}$$

Note that, $\mathbf{p}_1$ and $\mathbf{p}_2$ represent points in two different chessboard coordinate systems since the two cameras respectively capture each side of the chessboard. We further derive the rotation matrix and the translation matrix between the two chessboard coordinate systems. We use $\{\mathbf{R}_{\mathrm{chess}}, \mathbf{t}_{\mathrm{chess}}\}$ to represent them and have

$$\mathbf{p}_1 = \mathbf{R}_{\mathrm{chess}}\mathbf{p}_2 + \mathbf{t}_{\mathrm{chess}} \tag{5}$$

We have $\mathbf{R}_{\mathrm{chess}} = \mathrm{diag}\,(0, 0, -1)$ and $\mathbf{t}_{\mathrm{chess}} = (0, 0, -\mathrm{d})$, where d is the thickness of the chessboard. Note that, some cameras will capture images in a mirror mode. The rotation matrix should be adjusted based on real setting.

Therefore, given a point $\mathbf{p}_{\mathrm{depth}}$ in the depth camera coordinate system, we can obtain the $\mathbf{p}_{\mathrm{dms}}$ using Eq. (3), Eq. (4) and Eq. (5). We use $\mathbf{R}_{\mathrm{rot}}$ and $\mathbf{t}_{\mathrm{rot}}$ to represent the rotation and translation matrices between the depth and DMS cameras. It is easy to derive that

$$\mathbf{R}_{\mathrm{rot}} = \mathbf{R}_{\mathrm{dms}}\mathbf{R}_{\mathrm{chess}}\mathbf{R}_{\mathrm{depth}}^{-1}, \tag{6}$$

and

$$\mathbf{t}_{\mathrm{rot}} = -\mathbf{R}_{\mathrm{dms}}\mathbf{R}_{\mathrm{chess}}\mathbf{R}_{\mathrm{depth}}^{-1}\mathbf{t}_{\mathrm{depth}} + \mathbf{R}_{\mathrm{dms}}\mathbf{t}_{\mathrm{chess}} + \mathbf{t}_{\mathrm{dms}}. \tag{7}$$

We can use following equation for the conversion.

$$\mathbf{p}_{\mathrm{dms}} = \mathbf{R}_{\mathrm{rot}}\mathbf{p}_{\mathrm{depth}} + \mathbf{t}_{\mathrm{dms}}. \tag{8}$$

8.2. Implementation details of GazeDPTR

In this paper, we propose a GazeDPTR for gaze estimation. We also extend the GazeDPTR for gaze zone classification. We train the extended network in an end-to-end manner.

In detail, GazeDPTR contains two GazePTRs for feature extraction from original and normalized images. GazePTR is modified based on GazeTR [5]. We use ResNet18 to extract multi-level feature maps and obtain 4 different scale feature maps. Their scales are $64 \times 56 \times 56$, $128 \times 28 \times 28$, $256 \times 14 \times 14$, $512 \times 7 \times 7$. We use 1×1 convolution layers and global average pooling layers to convert them into 128D features. We denote these features as $\{f_i \in \mathbb{R}^{128}\}_{i=1,2,3,4}$. We use sup ^{n} and ^{o} to represent the feature is extracted from normalized or original images, e.g. f_1^o. Next, we use a 6-layer transformer to integrate these features for the final feature. We use one learnable token to aggregate $\{f_i^n\}$ for f_{final}^n. Two learnable tokens are used to aggregate $\{f_i^o\}$ since we need feature f_{final}^o for gaze estimation and visual feature f_{visual} for gaze zone classification. We use another 6-layer transformer to integrate f_{final}^n and f_{final}^o for f_{gaze}. We add a MLP to estimate gaze directions from f_{gaze}.

We project the estimated gaze into a tri-plane. Note that, we cut off the propagation of gradient in this operation layer since it drops gaze estimation accuracy but cannot improve gaze classification performance. We use a 2-layer transformer to extract positional feature f_{pos} where a deep transformer will vanish gradients. We also use a 6-layer transformer to integrate positional features and visual features for f_{zone}. We add a MLP to predict gaze zone from f_{zone}.

Regarding the loss function, we use L1 loss $\mathcal{L}_{gaze}$ for the gaze estimation task. Our method contains two ground truths $\mathbf{g}^o$ and $\mathbf{g}^n$. We define the function $\mathcal{L}_{gaze}^o(f)$ that means we set a MLP to estimate gaze from feature f and measure the L1 distance between the gaze and $\mathbf{g}^o$ for loss function. The same for $\mathcal{L}_{gaze}^n(f)$.

we require following feature should be gaze-related including 1) multi-level feature $\{f_i^n\}$ 2) intermediate features f_{final}^n and f_{final}^o 3) gaze feature f_{gaze}. The loss function can be represented as:

$$\mathcal{L}_1 = \sum_{i=1}^{4}\sum_{j\in\{o,n\}}\mathcal{L}_{gaze}^j(f_i^j) + \sum_{j\in\{o,n\}}\mathcal{L}_{gaze}^j(f_{final}^j) + \mathcal{L}_{gaze}^n(f_{gaze}) \tag{9}$$

We set cross entropy loss as the loss function for gaze zone classification. We also define the loss $\mathcal{L}_{zone}(f)$ that means we set a MLP to predict gaze zone from f and mea-

Table 7. We define nine zones for gaze zone classification and show the average precision (%) on each zone.

Visual Feature	Positional Feature	Left-side mirror	Rear-view mirror	Right-side mirror	Central-control screen	Steering wheel	Handbrake	Dashboard	Left-side windshield	Right-side windshield	Avg
✓		96.5	47.0	68.4	83.4	92.5	79.9	72.5	89.3	69.9	79.4
	✓	97.3	39.1	79.5	63.7	87.6	54.4	45.2	88.9	65.6	75.3
✓	✓	96.9	42.1	75.7	77.1	92.0	86.9	77.5	89.7	75.4	81.8

Table 8. We define one additional class *None* to account for samples that do not fall within the nine zones. We respectively report the average precision *with* and *w/o* the *None* region.

Visual Feature	Positional Feature	AP *w/o None* region	AP *with None* region
✓		79.4%	78.1%
	✓	75.3%	75.8%
✓	✓	81.8%	80.0%

Table 9. We selected best results from a pair of original and normalized images. The performance shows the selected result significantly outperform each of images.

	Original images	Normalized images	Selected
Acc	7.44°	7.04°	5.72°

sure the cross entropy loss. The loss function for gaze zone task is

$$\mathcal{L}_2 = \mathcal{L}_{zone}(f_{pos}) + \mathcal{L}_{zone}(f_{visual}) + \mathcal{L}_{zone}(f_{zone}) \quad (10)$$

We optimize the whole network using

$$\mathcal{L}_{GazeDPTR} = \mathcal{L}_1 + \mathcal{L}_2 \quad (11)$$

9. Additional Experiments

9.1. Setup of Gaze Zone Classification

Our paper extends gaze estimation for gaze zone classification. In this section, we provide details about the experimental setup.

During data collection, stickers are placed strategically within the vehicle. Based on the positions of these stickers, we divide the in-vehicle region into nine zones: left-side mirror, right-side mirror, rear-view mirror, steering wheel, left-side windshield, right-side windshield, central-control screen, handbrake, and dashboard. It is important to note that the dashboard encompasses not only the instrument cluster behind the steering wheel but also the air conditioning panel. Additionally, we introduce an extra region *None* to account for points that do not fall within the specified nine zones. The performance of this additional class is not included in the average performance calculation.

The average precision (AP) of each classes is shown in the Fig. 7. GazeDPTR integrates two features and show better average AP. We also show the average performance with *None* region in Tab. 8 for reference.
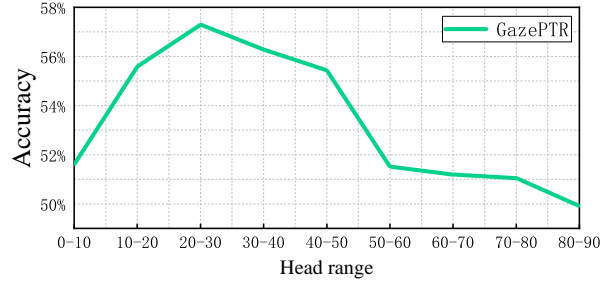

Figure 8. We count the improvement ratio in each head range. A larger ratio means more samples have performance improvement due to normalization. The result demonstrates that the large head range usually has relatively low improvement ratio.
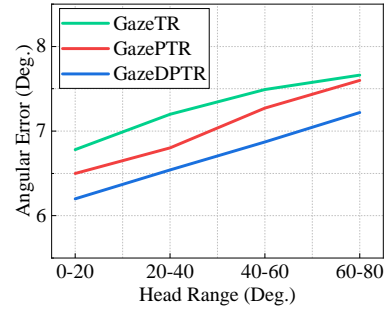

Figure 9. We count the average angular error in different head ranges. GazePTR estimates gaze from normalized images while GazeDPTR uses both normalized and original images for gaze estimation. It is interesting that GazeDPTR has larger performance improvement in a large head range than GazePTR. Combining with the result in Fig. 8, the reason may be the relatively low improvement ratio in the large head range.

9.2. Analysis on Normalized and Original Images

Our work uses both original and normalized images for gaze estimation. Our hypothesis is that the conbination of two images can provide additional insights beyond what each offers. In this section, we show experimental result to validate our hypothesis. We initially conducted an oracle baseline by separately training GazePTR on both the original and normalized datasets and selecting the best result from each image pair. The result is shown in Tab. 9. The selected performance demonstrated a remarkable improvement, achieving 5.72°, which significantly surpasses the performances in both the original and normalized images.

To gain a more nuanced understanding of the improvement across different head pose ranges, we calculated the

Figure 10. We show the normalization image obtained from Zhang *et al*. [38] and ours. We also visual the original images which are directly cropped from scene images. Zhang *et al*. rotate images based on the x-axis of head. It sometimes produce unstable result in extreme head pose, *e.g*., the second column. We modify their method and cancel such rotation. Our method has better performance which is shown in our manuscript.

improvement ratio. A particular sample is considered improved if the performance of the normalized image surpasses that of the original image. The results are visualized in Fig. 8, where images that failed in the large head pose range typically exhibit a relatively low improvement ratio. Additionally, the angular error across different head poses is depicted in Fig. 9, underscoring the larger performance improvement in a significant head pose range for GazeDPTR. These findings provide valuable insights into the advantages of our proposed method.

9.3. Visualization of Normalization Images

We show the images of different normalization methods and original images in Fig. 10. Zhang *et al*. [38] rotate images based on the x-axis of head. It sometimes produce unstable result in extreme head pose, *e.g*., the second column in Fig. 10. We modify their method and cancel such rotation. Our method has better performance which is shown in our manuscript.

References

[1] Abdul Rafey Aftab. Multimodal driver interaction with gesture, gaze and speech. In *2019 International Conference on Multimodal Interaction*, pages 487–492, 2019. 1

[2] Aya Ataya, Won Kim, Ahmed Elsharkawy, and SeungJun Kim. Gaze-head input: Examining potential interaction with immediate experience sampling in an autonomous vehicle. *Applied Sciences*, 10(24), 2020. 1

[3] Yiwei Bao, Yunfei Liu, Haofei Wang, and Feng Lu. Generalizing gaze estimation with rotation consistency. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4207–4216, 2022. 2

[4] Xin Cai, Jiabei Zeng, Shiguang Shan, and Xilin Chen. Source-free adaptive gaze estimation by uncertainty reduction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 22035–22045, 2023. 2

[5] Yihua Cheng and Feng Lu. Gaze estimation using transformer. *ICPR*, 2022. 3, 7, 1

[6] Yihua Cheng and Feng Lu. Dvgaze: Dual-view gaze estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 20632–20641, 2023. 3, 6

[7] Yihua Cheng, Feng Lu, and Xucong Zhang. Appearance-based gaze estimation via evaluation-guided asymmetric regression. In *The European Conference on Computer Vision*, 2018. 2

[8] Yihua Cheng, Shiyao Huang, Fei Wang, Chen Qian, and Feng Lu. A coarse-to-fine adaptive network for appearance-based gaze estimation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2020. 5

[9] Yihua Cheng, Xucong Zhang, Feng Lu, and Yoichi Sato. Gaze estimation by exploring two-eye asymmetry. *IEEE Transactions on Image Processing*, 29:5259–5272, 2020. 2

[10] Yihua Cheng, Haofei Wang, Yiwei Bao, and Feng Lu. Appearance-based gaze estimation with deep learning: A review and benchmark. *arXiv preprint arXiv:2104.12668*, 2021. 2, 4, 7

[11] Yihua Cheng, Yiwei Bao, and Feng Lu. Puregaze: Purifying gaze feature for generalizable gaze estimation. *AAAI*, 2022. 2, 3

[12] In-Ho Choi, Sung Kyung Hong, and Yong-Guk Kim. Real-time categorization of driver's gaze zone using the deep learning techniques. In *2016 International conference on big data and smart computing (BigComp)*, pages 143–148. IEEE, 2016. 2

[13] Tobias Fischer, Hyung Jin Chang, and Yiannis Demiris. Rt-gene: Real-time eye gaze estimation in natural environments. In *The European Conference on Computer Vision*, 2018. 2

[14] Lex Fridman, Philipp Langhans, Joonbum Lee, and Bryan Reimer. Driver gaze region estimation without use of eye movement. *IEEE Intelligent Systems*, 31(3):49–56, 2016. 2

[15] Kenneth Alberto Funes Mora, Florent Monay, and Jean-Marc Odobez. Eyediap: A database for the development and evaluation of gaze estimation algorithms from rgb and rgb-d cameras. In *Proceedings of the ACM Symposium on Eye Tracking Research & Applications*, 2014. 2

[16] Shreya Ghosh, Abhinav Dhall, Garima Sharma, Sarthak Gupta, and Nicu Sebe. Speak2label: Using domain knowledge for creating a large scale driver gaze zone estimation dataset. In *The IEEE/CVF International Conference on Computer Vision (ICCV) Workshops*, pages 2896–2905, 2021. 1, 2, 7

[17] Amie C Hayley, Brook Shiferaw, Blair Aitken, Frederick Vinckenbosch, Timothy L Brown, and Luke A Downey. Driver monitoring systems (dms): The future of impaired driving management? *Traffic injury prevention*, 22(4):313–317, 2021. 1

[18] Sumit Jha and Carlos Busso. Probabilistic estimation of the gaze region of the driver using dense classification. In *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, pages 697–702. IEEE, 2018. 1, 2

[19] Isaac Kasahara, Simon Stent, and Hyun Soo Park. Look both ways: Self-supervising driver gaze estimation and road scene saliency. In *European Conference on Computer Vision*, pages 126–142. Springer, 2022. 2

[20] Petr Kellnhofer, Adria Recasens, Simon Stent, Wojciech Matusik, and Antonio Torralba. Gaze360: Physically unconstrained gaze estimation in the wild. In *The IEEE International Conference on Computer Vision*, 2019. 1, 2, 7

[21] Muhammad Qasim Khan and Sukhan Lee. A comprehensive survey of driving monitoring and assistance systems. *Sensors*, 19(11), 2019. 1

[22] Muhammad Qasim Khan and Sukhan Lee. Gaze and eye tracking: Techniques and applications in adas. *Sensors*, 19 (24), 2019. 1

[23] Sung Joo Lee, Jaeik Jo, Ho Gi Jung, Kang Ryoung Park, and Jaihie Kim. Real-time gaze estimator based on driver's head orientation for forward collision warning system. *IEEE Transactions on Intelligent Transportation Systems*, 12(1): 254–267, 2011. 2

[24] Vincent Lepetit, Francesc Moreno-Noguer, and Pascal Fua. Epnp: An accurate o (n) solution to the pnp problem. *International journal of computer vision*, 81(2):155–166, 2009. 4

[25] Prajval Kumar Murali, Mohsen Kaboli, and Ravinder Dahiya. Intelligent in-vehicle interaction technologies. *Advanced Intelligent Systems*, 4(2):2100122, 2022. 1

[26] Seonwook Park, Adrian Spurr, and Otmar Hilliges. Deep pictorial gaze estimation. In *The European Conference on Computer Vision*, 2018. 2

[27] Seonwook Park, Emre Aksan, Xucong Zhang, and Otmar Hilliges. Towards end-to-end video-based eye-tracking. In *The European Conference on Computer Vision*, pages 747–763. Springer, 2020. 2

[28] Akshay Rangesh, Bowen Zhang, and Mohan M Trivedi. Driver gaze estimation in the real world: Overcoming the eyeglass challenge. In *The IEEE Intelligent Vehicles Symposium (IV)*, pages 1054–1059. IEEE, 2020. 2

[29] Yusuke Sugano, Yasuyuki Matsushita, and Yoichi Sato. Learning-by-synthesis for appearance-based 3d gaze estimation. In *The IEEE Conference on Computer Vision and Pattern Recognition*, 2014. 5

[30] Ashish Tawari, Kuo Hao Chen, and Mohan M Trivedi. Where is the driver looking: Analysis of head, eye and iris for robust gaze zone estimation. In *17th International IEEE conference on intelligent transportation systems (ITSC)*, pages 988–994. IEEE, 2014. 2

[31] Sourabh Vora, Akshay Rangesh, and Mohan Manubhai Trivedi. Driver gaze zone estimation using convolutional neural networks: A general framework and ablative analysis. *IEEE Transactions on Intelligent Vehicles*, 3(3):254–265, 2018. 1, 2

[32] Yafei Wang, Guoliang Yuan, Zetian Mi, Jinjia Peng, Xueyan Ding, Zheng Liang, and Xianping Fu. Continuous driver's gaze zone estimation using rgb-d camera. *Sensors*, 19(6): 1287, 2019. 1, 2

[33] Yaoming Wang, Yangzhou Jiang, Jin Li, Bingbing Ni, Wenrui Dai, Chenglin Li, Hongkai Xiong, and Teng Li. Contrastive regression for domain adaptation on gaze estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 19376–19385, 2022. 2

[34] Guoliang Yuan, Yafei Wang, Huizhu Yan, and Xianping Fu. Self-calibrated driver gaze estimation via gaze pattern learning. *Knowledge-Based Systems*, 235:107630, 2022. 2

[35] Mingfang Zhang, Yunfei Liu, and Feng Lu. Gazeonce: Real-time multi-person gaze estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4197–4206, 2022. 2, 6

[36] Xucong Zhang, Yusuke Sugano, Mario Fritz, and Andreas Bulling. Appearance-based gaze estimation in the wild. In *The IEEE Conference on Computer Vision and Pattern Recognition*, 2015. 2, 4

[37] Xucong Zhang, Yusuke Sugano, Mario Fritz, and Andreas Bulling. It's written all over your face: Full-face appearance-based gaze estimation. In *The IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 2299–2308, 2017. 2, 7

[38] Xucong Zhang, Yusuke Sugano, and Andreas Bulling. Revisiting data normalization for appearance-based gaze estimation. In *Proceedings of the ACM Symposium on Eye Tracking Research & Applications*, 2018. 5, 8, 3

[39] Xucong Zhang, Yusuke Sugano, Mario Fritz, and Andreas Bulling. Mpiigaze: Real-world dataset and deep appearance-based gaze estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(1):162–175, 2019. 1, 2

[40] Xucong Zhang, Seonwook Park, Thabo Beeler, Derek Bradley, Siyu Tang, and Otmar Hilliges. Eth-xgaze: A large scale dataset for gaze estimation under extreme head pose and gaze variation. In *The European Conference on Computer Vision*, 2020. 1, 2, 4, 7