

DriverGaze360: OmniDirectional Driver Attention with Object-Level Guidance

Shreedhar Govil¹

Didier Stricker^{1,2}

Jason Rambach¹

¹German Research Center for Artificial Intelligence (DFKI)

²Rhineland-Palatinate Technical University (RPTU)

shreedhar.govil@dfki.de, didier.stricker@dfki.de, jason.rambach@dfki.de

Abstract


Predicting driver attention is a critical problem for developing explainable autonomous driving systems and understanding driver behavior in mixed human-autonomous vehicle traffic scenarios. Although significant progress has been made through large-scale driver attention datasets and deep learning architectures, existing works are constrained by narrow frontal field-of-view and limited driving diversity. Consequently, they fail to capture the full spatial context of driving environments, especially during lane changes, turns, and interactions involving peripheral objects such as pedestrians or cyclists. In this paper, we introduce DriverGaze360, a large-scale 360° field of view driver attention dataset, containing ~1 million gaze-labeled frames collected from 19 human drivers, enabling comprehensive omnidirectional modeling of driver gaze behavior. Moreover, our panoramic attention prediction approach, DriverGaze360-Net, jointly learns attention maps and attended objects by employing an auxiliary semantic segmentation head. This improves spatial awareness and attention prediction across wide panoramic inputs. Extensive experiments demonstrate that DriverGaze360-Net achieves state-of-the-art attention prediction performance on multiple metrics on panoramic driving images. Dataset and method available at <https://dfki-av.github.io/drivergaze360>.

1. Introduction

Safe driving requires awareness of surroundings, constant monitoring of road and traffic, and alertness to react to unexpected events [21]. Advanced driving assistance systems (ADAS) such as forward collision warning and lane departure warning aim to prevent the consequences of distracted driving. The use of such technologies is promising but the overall effects are still unknown [2, 5]. Hence, it becomes imperative to build ADAS systems that can predict where a driver looks at during driving and alert the driver if needed [16, 17].


Predicting driver attention also plays an important role for

Current Driver Attention Datasets



Narrow Frontal Attention Maps

DriverGaze360



Complete 360° view Attention Maps

Figure 1. Existing methods predict driver attention only within a narrow frontal field of view, limiting understanding of gaze behavior. DriverGaze360 captures the full 360° field of view, enabling analysis of comprehensive gaze dynamics. In this example, the driver uses the mirrors to examine the rear-left while turning—a behavior not modeled in prior datasets.

autonomous vehicle development [8, 35, 37]. Gaze behavior encodes critical cues for decision-making and reaction timing. Modeling where and what a driver is looking at enables the design of explainable-AI systems that can anticipate human intent in mixed human-AI traffic scenarios [39].

Over the past decade, several large-scale datasets [11, 18, 32, 36] and learning-based models [1, 6, 12, 22, 39] have advanced the study of driver attention. However, prior work exclusively employs narrow, forward-facing cameras that capture only a fraction of the scene. This limited field of view restricts modeling of attention during complex maneuvers such as lane changes, turns, and interactions with lateral or rear entities like pedestrians or cyclists. Moreover, most existing datasets focus on brief, safety-critical events, offering limited insight into continuous, multi-directional gaze dynamics. The absence of omnidirectional data restricts progress toward holistic driver attention understanding.

To address this gap, we introduce DriverGaze360, a large-scale 360° field-of-view driver attention dataset that enables comprehensive omnidirectional modeling of gaze behavior. Collected in a controlled simulation environment, DriverGaze360 provides fine-grained control over peripheral objects, lane changes, and participant behavior. As shown in Figure 1, our dataset captures realistic gaze patterns such as rear-view monitoring—patterns that are unseen in existing work.

Simultaneously, we introduce DriverGaze360-Net, a vision transformer (ViT) based attention prediction model that jointly learns attention maps and the semantic categories of attended objects (vehicles, pedestrians, cyclists, and others) through an auxiliary segmentation head. This joint formulation improves spatial awareness and robustness when predicting sparse panoramic attention distributions. Extensive experiments demonstrate that DriverGaze360-Net achieves state-of-the-art performance across multiple saliency and segmentation metrics on panoramic data and generalizes effectively to real-world driving datasets.

We summarize the contributions of this paper as follows:

1. **DriverGaze360:** The first large-scale omnidirectional driver attention dataset collected from eye tracking of human drivers in a simulated environment, spanning diverse and safety-critical scenarios including lane changes, turns, and peripheral agent interactions.
2. **DriverGaze360-Net:** a transformer-based architecture that jointly predicts attention maps and attended objects through an auxiliary semantic segmentation head, enhancing spatial awareness and outperforming the current state-of-the-art models.
3. **Attended Object Extraction:** a fixation-semantic fusion pipeline that maps gaze distributions to object instances, yielding object-level attention annotations used as supervision for DriverGaze360-Net.

2. Related Work

2.1. Driver Attention Datasets

Over the last few years, there have been several works that build large scale datasets for capturing driver attention in var-

ious scenarios, thereby advancing understanding of human visual behavior in driving.

The DR(eye)VE [32] dataset was among the first to align captured gaze behavior with visual driving scenes, providing a foundation for learning-based attention prediction. LBW [18] expanded data collection to a broader driver pool under real-world conditions, while IVGaze [7] focused on in-vehicle gaze capture, but lacked corresponding driving footage. The BDD-A dataset by Xia et al. [36] focused on safety-critical events like emergency braking and traffic congestion. DADA-2000 [11] extended this to traffic accidents, including 2,000 videos and corresponding gaze data, thereby shifting the focus towards safety-critical and accident scenarios, offering large-scale annotated clips that expose the temporal dynamics of driver awareness under high-risk events.

Although these datasets have been instrumental, their scope remains limited. Most employ forward-facing cameras, representing only a fraction of the driver’s visual field. Consequently, they under-represent attention behaviors associated with lateral or rear spatial contexts—such as lane changes, merging, or monitoring pedestrians in blind spots. To address these constraints, we present DriverGaze360, a human driver-synthetic scene 360° driver attention dataset designed to capture full-surround gaze behavior. Unlike previous frontal-view datasets, DriverGaze360 records panoramic driving environments rendered in simulation, enabling the modeling of attention beyond the windshield and full control over the captured driving scenarios. Table 1 summarizes how DriverGaze360 differs from existing datasets.

2.2. Driver Attention Prediction

Driver attention prediction has advanced alongside the rise of large-scale driver-attention datasets. Early work such as DR(eye)VE [32] employed a multi-sensor CNN pipeline using RGB, semantic segmentation, and optical flow to approximate human gaze. BDD-A [36] leveraged AlexNet [23] features with ConvLSTMs to capture short-term temporal cues from front-facing video. SCAFFNet [12] is a method of simultaneous processing and fusion of RGB and semantic images. SAGE [31] treats all road users as a single foreground class, adding it to saliency ground truth, and performing foreground-background classification. ASIAFNet [25] combines short-temporal motion features with object-level attention estimation, using bounding box detection and binary classification. While MEDIRL [1] used maximum entropy depth inverse reinforcement learning. More recent approaches explore richer context: FBLNet [6] introduces a feedback loop to simulate driver experience and uses dual CNN-Transformer encoders; SCOUT+ [22] encodes bird’s-eye-view road geometry via a map encoder and couples it with image embeddings to improve driver attention prediction at safety-critical scenarios and intersections.

Table 1. Comparison of DriverGaze360 with existing driver attention datasets.

Dataset	360° FoV	# Hours	Scenarios	# Subjects	Data Collection
DR(eye)VE [32]	×	6	Regular Driving	8	Real driving
LBW [18]	×	7	Regular Driving	28	Real driving
BDD-A [36]	×	4	Busy Intersections, Emergency Breaking	1,228	Watching videos
DADA-2000 [11]	×	6	Driving Accidents	20	Watching videos
DriverGaze360 (ours)	✓	9	Regular Driving, Critical Situations	19	Simulated driving

However, these methods operate on narrow field-of-view forward imagery and thus fail to represent the full 360° attention field that real drivers deploy. We show their performance degrades on panoramic inputs, limiting applicability to omnidirectional perception. In contrast, DriverGaze360-Net pairs each panoramic frame with object-level semantic labels (e.g., vehicles, pedestrians, traffic signals) to identify attended entities, yielding stronger saliency performance and improved real-world generalization.

3. Building the DriverGaze360 Dataset

In this section, we introduce our setup for gaze data collection of human drivers in simulation environments. Our system includes a driving simulator, eye tracking setup, scenario player, and a post processing step to build attention maps. Leveraging this system with 19 driver subjects, we build the first 360° driver attention dataset. The dataset spans a diverse range of scenarios, facilitating and advancing future research on omnidirectional attention prediction.

3.1. Data Collection Setup

We use the CARLA simulator [9] due to its open-source support, reproducibility, and broad use in autonomous-driving research. CARLA provides configurable camera, LiDAR, and control interfaces, and supports flexible scenario generation. We employ both its native scenario tools and SCENIC [13] to produce varied traffic, agent behaviors, and weather conditions.

To approximate a naturalistic driving field of view, the setup employs three large front displays and two picture-in-picture displays serving as rear-view mirrors with 72° FoV each, forming a contiguous 360° FoV. This configuration enables acquisition of wide-angle gaze patterns, including peripheral and backward attention shifts. The setup is shown in Figure 2 (left).

We capture the eye gaze and the driver’s egocentric perspective through the Pupil Core [15] eye-tracking glasses. The device operates at 120 Hz for eye-tracking and provides a synchronized egocentric video at 30 Hz. Participants operate a cockpit system with steering, throttle, and brake, plus gear control. A heads-up display shows speed and current gear. AprilTags [30] rendered in the simulated scene enable calibration between eye-tracker coordinates and simulator

space.

3.2. Traffic Scenarios

To mitigate sim-to-real bias and ensure broad attention variability, we design scenarios spanning routine, goal-directed, and safety-critical driving. All scenarios are parameterized by weather, time-of-day, road type, traffic density, and pedestrian behavior. Parameters are randomized at episode start. We consider three scenario classes:

- Unscripted navigation.** Participants drive freely in mixed urban-suburban environments. They choose their own routes and maneuvers.
- Goal-directed navigation.** Participants follow audio-guided navigation through a series of checkpoints such as road construction, pedestrian crosswalks, multilane roundabouts, highway merges, etc.
- Safety-critical events.** We inject five common near-miss situations with randomized timing, actor types, and traffic density:
 - Highway emergency braking: sudden lead-vehicle deceleration with blocked adjacent lanes.
 - Highway merging: entering dense traffic from the right lane.
 - Urban pedestrian crossing: partially occluded pedestrian enters roadway.
 - Signalized left turns with oncoming-vehicles.
 - Highway cut-in with short-headway insertion by front vehicle.

Across all categories, scene layouts, actor trajectories, and distractors are varied to increase behavioral coverage while maintaining reproducibility. Figure 3 shows data samples with varying weather and traffic conditions. More information regarding scenarios is provided in Section 7 where Figure 11 illustrates the safety-critical scenarios that were implemented.

3.3. Fixation Target Calibration

We align eye-tracker fixation points from the eye tracker coordinate frame to the CARLA simulator using AprilTags [30], shown in Figure 2 (left). Each gaze point is mapped to the CARLA image plane via a homography. Similarly to prior-work [32], we build a fixation map by aggregating gaze within a 30-frame temporal window centered at time t : gaze




Figure 2. Experimental Setup. (Left) Egocentric perspective during a data collection with rainy weather and a pedestrian. (Right) Resulting attention maps with a 360° FoV.



Figure 3. DriverGaze360 samples with different weather and traffic conditions.




Figure 4. Fixation Distribution for DriverGaze360

points from neighboring frames are projected into frame t , converted to 2-D Gaussians with fixed spatial variance, and finally normalized to form a probability map, Figure 2 (right). More details are provided in Section 8.

3.4. Participant Data Collection

We recruited 21 licensed drivers (age 21–40) with at least three years of driving experience. Participants received a 5 min on-boarding session to familiarize with the driving simulator and eye-tracking glasses. During data collection, participants were instructed to follow traffic regulations. Two subjects were excluded due to motion sickness or difficulty controlling the simulator.

Each session consisted of multiple driving episodes under varied conditions, including daylight, nighttime, wind, and rain, using standard CARLA weather presets. Simulator video streams were sampled at 30 FPS to match the eye-tracker acquisition rate.

3.5. Dataset Properties

We introduce DriverGaze360, the first large-scale 360° driver attention dataset that provides dense gaze annotations, human driving behavior, and multi-sensor outputs. By leveraging CARLA’s simulation and scenario replay capabilities, the dataset supports both traditional RGB-based gaze prediction and multi-sensor fusion approaches.

Dataset statistics: DriverGaze360 contains approximately 9 h of driving footage with 1 M gaze-annotated frames (obtained from 5 cameras with 1 M frame each), collected from 19 participants, with per-driver contributions ranging from 18K to 126K frames. Data were recorded across three scenario types under diverse environmental conditions (day/night, rain, wind): ~ 80 minutes of unscripted navigation, ~ 85 minutes of safety-critical events, and ~ 370 minutes of goal-directed navigation. Rear-view fixations account for 6% of all gaze samples, aligning with reported real-world rates of 5-10% [3, 4]. A visualization of the gaze distribution is shown in Figure 4. Our dataset is sampled at 30 Hz and for each timestep it outputs five synchronized RGB frames at 1280×720 resolution.

Comparison to existing datasets: As evident in Table 1,

DriverGaze360 provides a full 360° field of view in a simulated environment, enabling comprehensive modeling of driver attention under both routine and challenging scenarios. Even though our data is collected within a simulated environment it involves real human driving, unlike BDDA [36] or DADA-2000 [11], which were collected by participants watching driving videos without direct engagement in driving. Moreover, we provide 9 hours of driving footage, more than any previous dataset, while covering both regular driving and critical situations.

Data Splits. Data were collected in standard CARLA maps referred to as Towns, namely Towns {1, 2, 3, 4, 5, 6, 7, 10, and 11}. We follow a map-based split strategy: Towns {2, 3, 4, 7, 10, 11} for training and Towns {1, 5, 6} for validation. This enforces non-overlapping geographic layouts while maintaining comparable distributions of urban, suburban, and highway scenarios. The resulting dataset is balanced across the splits with respect to duration.

4. DriverGaze360-Net

Omnidirectional scene capturing introduces new settings and challenges to driver attention which we systematically explore in this section. We introduce our attention prediction network, DriverGaze-Net and an attended object mask generation algorithm.

4.1. Network Architecture

The proposed DriverGaze360-Net architecture, illustrated in Figure 5, comprises three main components: a ViT-based [10] scene encoder for spatio-temporal feature extraction, an attention decoder for driver attention prediction, and a segmentation decoder for attended objects.

We adopt the Video Swin Transformer (VST) [27] as the backbone to extract multi-scale spatio-temporal features from sequences of RGB frames through the four encoder stages. VST was chosen because of its hierarchical architecture and shifted window-based self-attention mechanism [26]. By performing attention calculations only within local, non-overlapping windows and shifting these windows between layers, it reduces computational complexity from quadratic (as with standard ViT) to linear with respect to image size.

The spatio-temporal features are fused through a shared convolutional decoder that progressively up-samples and aggregates encoder outputs using Conv3D layers and ReLU activations which enables weight sharing between downstream decoders. The attention decoder refines the shared representation using additional convolutional layers, consisting of Conv3D and ReLU activation blocks, with a final upsampling layer, to produce dense 360° gaze probability maps, while the attended object decoder mirrors this structure but includes a final classification layer to predict attended object categories and per-pixel scores, allowing object-level

semantic reasoning aligned with driver attention. We define “attended objects” as active traffic participants (vehicles, pedestrians, cyclists), traffic signs and traffic lights that are perceived by the driver.

The network operates on sequences of concatenated RGB frames from the past T time steps, resulting in an input shape of $T \times 3 \times H \times W$. Features are first extracted via the VST encoder and then fed to the downstream decoders. The outputs of the network are an omnidirectional gaze attention map from the attention decoder and a segmentation map of attended objects from the attended object decoder.

4.2. Attended Object Segmentation

We jointly train a semantic segmentation head to predict attended objects, rather than treating attention prediction and semantic segmentation as independent tasks. Prior saliency networks such as [12, 32] typically utilize full-scene semantic segmentation as input without explicitly distinguishing which objects are relevant to the driving context. This limitation becomes more pronounced in panoramic data, such as our DriverGaze360 dataset, where attention signals are inherently sparse and localized. Under such conditions, it is essential to focus the supervision on the attention-relevant objects rather than the entire scene.

To supervise object-aware attention, we generate new masks that retain only objects attended by the driver. Since most objects in 360° scenes are unseen at a given moment, full semantic supervision is not optimal, as we show in our ablation study Section 5.3. We first isolate road-user instances $R_{road} = \{\text{vehicles, pedestrians, cyclists, traffic signs, traffic lights}\}$. The gaze map S_{sal} is binarized to obtain a salient region mask, which is intersected with each instance. Instances with non-empty intersection are labeled as attended. We then construct an attended-object map S_{obj} in which only pixels belonging to attended instances retain their semantic IDs; all others are set to background. The procedure is provided in Algorithm 1, where $\odot$ is element-wise multiplication and $1[\cdot]$ is an indicator function.

4.3. Metrics

Various metrics can be used to compare attention maps. We employ four widely-used saliency metrics [24]: Kullback-Leibler Divergence (KLD) that measure distance between two probability distributions, Correlation Coefficient (CC) which calculates the linear relationship between the ground truth and prediction, Similarity Index (SIM) which indicates the similarity, and Normalized Scanpath Saliency (NSS) which measures the difference, formulated in Table 2. For semantic segmentation, we use two common metrics, Dice coefficient and Intersection-over-Union (IoU), as defined in Table 3.




Figure 5. DriveGaze360-Net Architecture. We introduce a novel attended object prediction head that jointly learns an attention map and attended object segmentation. This addition improves the performance under wide panoramic images with highly sparse attention maps.

Algorithm 1: Attended Object Extraction

Input: Attention map S_{sal} , Instance segmentation I_{inst} , τ threshold, Road users R_{road}

Output: Semantic mask of attended objects S_{obj}

1 Extract road-user instances:

$$I_{road} \leftarrow \{I_i \in I_{inst} \mid \text{class}(I_i) \in R_{road}\};$$

2 Binarize attention map: $\hat{S}_{sal} \leftarrow \mathbb{1}[S_{sal} > \tau]$;

3 Compute overlap: $M_{sal} \leftarrow \hat{S}_{sal} \odot I_{road}$;

4 Identify attended instances:

$$I_{sal} \leftarrow \{I_i \in I_{road} \mid M_{sal} \cap I_i \neq \emptyset\};$$

5 Generate attended-object segmentation for all pixels p :

$$S_{obj}(p) = \begin{cases} \text{class_id}(I_i), & \text{if } p \in I_i, I_i \in I_{sal}; \\ 0, & \text{otherwise} \end{cases}$$

6 **return** S_{obj} ;

Metric	Formulations	Range
KLD ↓	$\sum_i P_X(i) \log(\frac{P_X(i)}{\bar{P}_Y(i)+\epsilon} + \epsilon)$	$[0, \infty)$
CC ↑	$\frac{\text{cov}(P_X, P_Y)}{\sigma(P_X)\sigma(P_Y)}$	$[0, 1]$
NSS ↑	$\frac{1}{N \neq 0} \sum_i \bar{P}_Y(i) \cdot F(i)$	$[0, \infty)$
SIM ↑	$\sum_i \min(P_X(i), P_Y(i))$	$[0, 1]$

Table 2. Metrics for attention map comparison. The symbols ↓ and ↑ indicate that lower or higher values are preferred, respectively. P_X and P_Y represent the ground-truth and predicted probability distributions of attention maps, respectively. $\bar{P}$ is the normalized probability distribution, F denotes the fixation map for an image, $\text{cov}(P_X, P_Y)$ is the covariance, and $\sigma(\cdot)$ denotes the standard deviation. The summation index i runs over image pixels, and ϵ is a small constant for numerical stability.

Metric	Formulations	Range
Dice ↑	$2 \frac{ X \cap Y }{ X + Y }$	$[0, 1]$
IoU ↑	$\frac{ X \cap Y }{ X \cup Y }$	$[0, 1]$

Table 3. Evaluation metrics for semantic segmentation. X and Y denote the sets of ground-truth and predicted pixels for a given class.

4.4. Loss Function

Our model jointly optimizes attention prediction and attended object segmentation. For attention prediction, we use $\mathcal{L}_{sal}$ as defined in Equation (1), which combines the KL Divergence with the negative correlation coefficient between the ground-truth and predicted attention maps, X_{sal} and Y_{sal} :

$$\mathcal{L}_{sal}(X_{sal}, Y_{sal}) = KLD(P_{X_{sal}}, P_{Y_{sal}}) - CC(X_{sal}, Y_{sal}) \quad (1)$$

For attended object segmentation, the loss $\mathcal{L}_{seg}$ as defined in Equation (2) maximizes the Dice and IoU scores while minimizing the cross-entropy $\mathcal{L}_{CE} = -\sum_{c=1}^N y_c \log p_c$ between the ground-truth X_{seg} and predicted Y_{seg} segmentation maps:

$$\mathcal{L}_{seg}(X_{seg}, Y_{seg}) = -\text{Dice}(X_{seg}, Y_{seg}) - \text{IoU}(X_{seg}, Y_{seg}) + \mathcal{L}_{CE}(X_{seg}, Y_{seg}) \quad (2)$$

The overall training objective is a weighted combination of these two losses, defined as $\mathcal{L}$ in Equation (3) with weights

λ_{sal} and λ_{seg} . In our experiments, we set $\lambda_{sal} = \lambda_{seg} = 1$:

$$\mathcal{L}(X, Y) = \lambda_{sal} \mathcal{L}_{sal}(X_{sal}, Y_{sal}) + \lambda_{seg} \mathcal{L}_{seg}(X_{seg}, Y_{seg}) \quad (3)$$

5. Evaluation

In this section, we outline the experimental setup, present quantitative and qualitative results, and provide an in-depth analysis. We further assess real-world generalization using an external dataset and demonstrate the benefit of object-level guidance for attention prediction in an ablation experiment.

5.1. Experimental Setup

Datasets. For experiments on DriverGaze360, the input is formed by horizontally concatenating the five RGB camera views and resizing to 1120×224 . Each training example consists of 16 uniformly sampled consecutive frames ($\sim 0.5s$), producing an input tensor of size $16 \times 3 \times 224 \times 1120$. Instance segmentation masks are extracted directly from CARLA. All evaluated methods are trained on the DriverGaze360 training set for fair comparison. Due to the absence of real panoramic driver attention datasets, we perform an experiment on the DADA-2000 [11] dataset with a narrow FoV to show that our method is widely applicable. For these evaluations, the images are resized to 224×224 , and YOLO-v11 [20] is used for instance segmentation.

Network. We train our network, DriverGaze360-Net, using loss from Eq. (3), with the AdamW optimizer [28] (betas: 0.9, 0.999; weight decay: 0.01) using a fixed learning rate of 1×10^{-4} and batch size of 4. We use the Swin-S backbone pretrained on Kinetics-400 [19]. Our model is trained for 20 epochs with early stopping on a single NVIDIA H100 GPU, requiring approximately 24 hours.

5.2. Comparison with SOTA Methods

In order to illustrate the performance of our method, we compared it with five other SOTA methods, among which are Dr(eye)VE [32], BDDA [36], DADANet [12], ViNet++ [14], and FBLNet [6].

Quantitative Comparison. We show the result on DriverGaze360 in Table 4. We can observe that our method outperforms all other methods. Compared with the second-best result on KLD, SIM, CC and NSS metrics, our model achieves 12.18%, 4.24%, 4.51%, 4.94% performance improvement, respectively, demonstrating our method’s advantage in wide panoramic inputs, which can also be seen qualitatively in Figure 6.

Table 5 presents the results for the real-world DADA-2000 [11] dataset, here our method outperforms all baselines and improves on the second best method in SIM, CC, NSS by 7.32%, 4.55%, 3.36% respectively, and is outperformed on KLD by just 0.48%. This shows that our method is not dataset-specific and is applicable to real-world data as well, which is also demonstrated qualitatively in Figure 7.

Model	KLD ↓	SIM ↑	CC ↑	NSS ↑
Dr(eye)VE [32]	1.293	0.340	0.613	5.452
BDDA [36]	2.566	0.218	0.400	2.772
DADANet [12]	1.269	0.476	0.618	5.598
ViNet++ [14]	1.251	0.442	0.611	2.772
FBLNet [6]	1.215	0.494	0.639	6.012
DriverGaze360-Net (ours)	1.067	0.515	0.667	6.309

Table 4. Results on DriverGaze360 Dataset.




Figure 6. Qualitative results on DriverGaze360 (nighttime).

Model	KLD ↓	SIM ↑	CC ↑	NSS ↑
Dr(eye)VE [32]	2.065	0.325	0.451	2.920
BDDA [36]	1.820	0.290	0.440	2.805
DADANet [12]	1.646	0.353	0.484	3.365
ViNet++ [14]	1.719	0.352	0.472	3.234
FBLNet [6]	1.818	0.369	0.480	3.305
DriverGaze360-Net (ours)	1.654	0.396	0.506	3.478

Table 5. Results on the DADA-2000 Dataset.




Figure 7. Qualitative results on DADA-2000 (motorbike accident).




Figure 8. Attention prediction and Attended Object segmentation in DriverGaze360

Heads	KLD ↓	SIM ↑	CC ↑	NSS ↑	Dice ↑	IoU ↑
Attention	1.127	0.510	0.654	6.158	-	-
+ ObjSeg.	<u>1.092</u>	<u>0.510</u>	<u>0.659</u>	<u>6.167</u>	<u>0.636</u>	<u>0.597</u>
+ AttObjSeg.	1.067	0.515	0.667	6.309	0.639	0.626

Table 6. Ablation results on attended-object supervision.

5.3. Effect of Attended Object Segmentation

Table 6 reports the impact of adding the attended-object segmentation head. The baseline predicts only attention maps. Adding a semantic object segmentation head (ObjSeg), i.e. segmenting all objects regardless of the attention, yields




Figure 9. Attention prediction and Attended Object segmentation in DADA-2000 during a motorbike accident.

consistent improvements, indicating that object-level cues help refine spatial attention. Introducing the attended-object segmentation head (AttObjSeg) further boosts all attention metrics, confirming that explicitly supervising the model to identify the viewer-relevant objects provides stronger guidance.

Relative to the attention-only baseline, AttObjSeg improves KLD, SIM, CC, and NSS by 5.32%, 1.11%, 2.08%, and 2.45%, respectively; compared to ObjSeg, it provides additional gains of 2.29%, 1.08%, 1.31%, and 2.30%. Qualitatively, the attended-object head successfully highlights task-critical regions (e.g., traffic lights, vehicles, and vulnerable road users), as shown in Figure 8 and Figure 9.

6. Conclusion

In this work, we present a comprehensive study on omnidirectional driver attention, encompassing a new dataset, predictive algorithm, and extensive experimental validation. We introduced DriverGaze360, the first large-scale 360° driver attention dataset with dense gaze annotations, enabling detailed modeling of human visual behavior in full-surround driving environments. Building on this, we proposed a transformer-based model with an auxiliary semantic segmentation head that jointly predicts attention maps and attended objects, improving both spatial coherence and interpretability. Our findings offer two key insights: 1) an omnidirectional dataset enables holistic modeling of driver awareness. 2) joint attended-object prediction significantly enhances panoramic attention estimation by guiding the model toward semantically meaningful regions. For future work, we plan to leverage DriverGaze360 to develop explainable AI driver models that integrate human attention cues into autonomous decision-making, advancing transparency and reliability in autonomous vehicle systems [33, 34, 39].

Acknowledgments

This work was partially funded by the European Union's Horizon Europe Research and Innovation Programme under Grant Agreement No. 101076360 (BERTHA) and by the German Federal Ministry of Research, Technology and Space under Grant Agreement No. 16IW24009 (COPPER).

The authors would like to express their sincere appreciation to Prateek Kumar Sharma, for his support with data collection and the implementation of driving scenarios. We also gratefully acknowledge Ruben Abad, Alex Levy, and Prof. Antonio M. López from the Computer Vision Center (CVC) for their methodological guidance and for providing the code used to implement the goal-directed navigation routes applied in collecting part of the dataset presented in this study. Finally, we sincerely thank all the participants who contributed to the dataset collection, as well as our colleagues at DFKI for their valuable feedback and support throughout this project.

References

- [1] Sonia Bae, Erfan Pakdamanian, Inki Kim, Lu Feng, Vicente Ordonez, and Laura Barnes. Medirl: Predicting the visual attention of drivers via maximum entropy deep inverse reinforcement learning. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 13158–13168, 2021. 2
- [2] Lyndel Bates, Marina Alexander, Margo Van Felius, John Seccombe, and Emma Bures. Final Report - What is known about distracted driving? - Australian Automobile Association — [aaa.asn.au](https://www.aaa.asn.au/library/final-report-what-is-known-about-distracted-driving/). <https://www.aaa.asn.au/library/final-report-what-is-known-about-distracted-driving/>, 2024. [Accessed 09-11-2025]. 1
- [3] Francesco Biondi, Praneet Sahoo, and Noor Jajo. The distraction potential of driving a partially automated vehicle through a construction zone. *Scientific Reports*, 15(1):8539, 2025. 4
- [4] Stewart Birrell and Mark Fowkes. Glance behaviours when using an in-vehicle smart driving aid: A real-world, on-road driving study. *Transportation Research Part F: Traffic Psychology and Behaviour*, 22:113–125, 2014. 4
- [5] Christopher D. D. Cabral, Jork C. J. Stapel, Riender Happee, and Joost C. F. de Winter. Redesigning today's driving automation toward adaptive backup control with context-based and invisible interfaces. *Human Factors*, 62(2):211–228, 2020. PMID: 31995390. 1
- [6] Yilong Chen, Zhixiong Nan, and Tao Xiang. Fblnet: Feedback loop network for driver attention prediction. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 13325–13334, 2023. 2, 7
- [7] Yihua Cheng, Yaning Zhu, Zongji Wang, Hongquan Hao, Yongwei Liu, Shiqing Cheng, Xi Wang, and Hyung Jin Chang. What do you see in vehicle? comprehensive vision solution for in-vehicle gaze estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1556–1565, 2024. 2
- [8] Ning Ding, Ce Zhang, and Azim Eskandarian. Saliendet: A saliency-based feature enhancement algorithm for object detection for autonomous driving. *IEEE Transactions on Intelligent Vehicles*, 9(1):2624–2635, 2024. 1
- [9] Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. Carla: An open urban driving simulator. In *Conference on robot learning*, pages 1–16. PMLR, 2017. 3
- [10] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021. 5
- [11] Jianwu Fang, Dingxin Yan, Jiahuan Qiao, Jianru Xue, He Wang, and Sen Li. Dada-2000: Can driving accident be predicted by driver attentionf analyzed by a benchmark. In *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*, pages 4303–4309, 2019. 2, 3, 5, 7
- [12] Jianwu Fang, Dingxin Yan, Jiahuan Qiao, Jianru Xue, and Hongkai Yu. Dada: Driver attention prediction in driving accident scenarios. *IEEE Transactions on Intelligent Transportation Systems*, 23(6):4959–4971, 2022. 2, 5, 7
- [13] Daniel J Fremont, Edward Kim, Tommaso Dreossi, Shromona Ghosh, Xiangyu Yue, Alberto L Sangiovanni-Vincentelli, and Sanjit A Seshia. Scenic: a language for scenario specification and data generation. *Machine Learning*, 112(10):3805–3849, 2023. 3, 1
- [14] Rohit Girmaji, Siddharth Jain, Bhav Beri, Sarthak Bansal, and Vineet Gandhi. Minimalistic video saliency prediction via efficient decoder & spatio temporal action cues. In *ICASSP 2025 - 2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2025. 7
- [15] Pupil Labs GmbH. Pupil Core - Open source eye tracking platform — [pupil-labs.com](https://pupil-labs.com/products/core). <https://pupil-labs.com/products/core>, 2020. [Accessed 09-11-2025]. 3
- [16] Feiyan Hu, Venkatesh G M, Noel E. O'Connor, Alan F. Smeaton, and Suzanne Little. Utilising visual attention cues for vehicle detection and tracking. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 5535–5542, 2021. 1
- [17] Muhammad Monjurul Karim, Yu Li, Ruwen Qin, and Zhaozheng Yin. A dynamic spatial-temporal attention network for early anticipation of traffic accidents. *IEEE Transactions on Intelligent Transportation Systems*, 23(7):9590–9600, 2022. 1
- [18] Isaac Kasahara, Simon Stent, and Hyun Soo Park. Look Both Ways: Self-supervising Driver Gaze Estimation and Road Scene Saliency. In *Computer Vision – ECCV 2022*, pages 126–142, Cham, 2022. Springer Nature Switzerland. 2, 3
- [19] Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, et al. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*, 2017. 7
- [20] Rahima Khanam and Muhammad Hussain. Yolov11: An

- overview of the key architectural enhancements. *arXiv preprint arXiv:2410.17725*, 2024. 7
- [21] Neale Kinnear and Alan Stevens. The Battle for Attention: Driver Distraction — docslib.org. <https://docslib.org/doc/13264384/the-battle-for-attention-driver-distraction>, 2015. [Accessed 09-11-2025]. 1
- [22] Iuliia Kotseruba and John K. Tsotsos. Scout+: Towards practical task-driven drivers’ gaze prediction. In *2024 IEEE Intelligent Vehicles Symposium (IV)*, pages 1927–1932, 2024. 2
- [23] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. ImageNet classification with deep convolutional neural networks. *Commun. ACM*, 60(6):84–90, 2017. 2
- [24] Matthias Kummerer, Thomas SA Wallis, and Matthias Bethge. Saliency benchmarking made easy: Separating models, maps and metrics. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 770–787, 2018. 5
- [25] Qiang Li, Chunsheng Liu, Faliang Chang, Shuang Li, Hui Liu, and Zehao Liu. Adaptive short-temporal induced aware fusion network for predicting attention regions like a driver. *IEEE Transactions on Intelligent Transportation Systems*, 23(10):18695–18706, 2022. 2
- [26] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022, 2021. 5
- [27] Ze Liu, Jia Ning, Yue Cao, Yixuan Wei, Zheng Zhang, Stephen Lin, and Han Hu. Video swin transformer. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3192–3201, 2022. 5
- [28] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 7
- [29] Daniel Martin, Ana Serrano, and Belen Masia. Panoramic convolutions for 360° single-image saliency prediction. In *CVPR Workshop on Computer Vision for Augmented and Virtual Reality*, 2020. 3
- [30] Edwin Olson. AprilTag: A robust and flexible visual fiducial system. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 3400–3407. IEEE, 2011. 3, 2
- [31] Anwesan Pal, Sayan Mondal, and Henrik I. Christensen. “looking at the right stuff” - guided semantic-gaze for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2
- [32] Andrea Palazzi, Davide Abati, simone Calderara, Francesco Solera, and Rita Cucchiara. Predicting the driver’s focus of attention: The dr(eye)ve project. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(7):1720–1733, 2019. 2, 3, 5, 7
- [33] Hao Shao, Yuxuan Hu, Letian Wang, Guanglu Song, Steven L. Waslander, Yu Liu, and Hongsheng Li. Lmdrive: Closed-loop end-to-end driving with large language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15120–15130, 2024. 8
- [34] Chonghao Sima, Katrin Renz, Kashyap Chitta, Li Chen, Hanxue Zhang, Chengen Xie, Jens Beißwenger, Ping Luo, Andreas Geiger, and Hongyang Li. Drivelm: Driving with graph visual question answering. In *European conference on computer vision*, pages 256–274. Springer, 2024. 8
- [35] Ke Wang, Sai Ma, Fan Ren, and Jianbo Lu. Sbas: Salient bundle adjustment for visual slam. *IEEE Transactions on Instrumentation and Measurement*, 70:1–9, 2021. 1
- [36] Ye Xia, Danqing Zhang, Jinkyu Kim, Ken Nakayama, Karl Zipser, and David Whitney. Predicting driver attention in critical situations. In *Asian conference on computer vision*, pages 658–674. Springer, 2018. 2, 3, 5, 7
- [37] Ye Xia, Jinkyu Kim, John Canny, Karl Zipser, Teresa Canas-Bajo, and David Whitney. Periphery-fovea multi-resolution driving model guided by human attention. In *2020 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1756–1764, 2020. 1
- [38] Heeseung Yun, Sehun Lee, and Gunhee Kim. Panoramic vision transformer for saliency detection in 360° videos. In *Computer Vision – ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXV*, page 422–439, Berlin, Heidelberg, 2022. Springer-Verlag. 3
- [39] Yuchen Zhou, Jiayu Tang, Xiaoyan Xiao, Yueyao Lin, Linkai Liu, Zipeng Guo, Hao Fei, Xiaobo Xia, and Chao Gou. Where, what, why: Towards explainable driver attention prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 2675–2685, 2025. 1, 2, 8