



UNIVERSITY OF LEEDS

This is a repository copy of *Driver-Net: Multi-Camera Fusion for Assessing Driver Take-Over Readiness in Automated Vehicles*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/229641/>

Version: Accepted Version

---

**Proceedings Paper:**

Rezaei, M. [orcid.org/0000-0003-3892-421X](https://orcid.org/0000-0003-3892-421X) and Azarmi, M. (Accepted: 2025) Driver-Net: Multi-Camera Fusion for Assessing Driver Take-Over Readiness in Automated Vehicles. In: IEEE Intelligent Vehicles Symposium Proceedings. 36th IEEE Intelligent Vehicles Symposium 2025, 22-25 Jun 2025, Napoca, Romania. IEEE (In Press)

---

This is an author produced version of a conference paper accepted for publication in IEEE Symposium on Intelligent Vehicle Proceedings made available under the terms of the Creative Commons Attribution License (CC-BY), which permits unrestricted use, distribution and reproduction in any medium, provided the original work is properly cited.

**Reuse**

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here:

<https://creativecommons.org/licenses/>

**Takedown**

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing [eprints@whiterose.ac.uk](mailto:eprints@whiterose.ac.uk) including the URL of the record and the reason for the withdrawal request.



[eprints@whiterose.ac.uk](mailto:eprints@whiterose.ac.uk)  
<https://eprints.whiterose.ac.uk/>

# Driver-Net: Multi-Camera Fusion for Assessing Driver Take-Over Readiness in Automated Vehicles\*

Mahdi Rezaei<sup>1,†</sup> and Mohsen Azarmi<sup>1</sup>

**Abstract**—Ensuring safe transition of control in automated vehicles requires an accurate and timely assessment of driver readiness. This paper introduces Driver-Net, a novel deep learning framework that fuses multi-camera inputs to estimate driver take-over readiness. Unlike conventional vision-based driver monitoring systems that focus on head pose or eye gaze, Driver-Net captures synchronised visual cues from the driver’s head, hands, and body posture through a triple-camera setup. The model integrates spatio-temporal data using a dual-path architecture, comprising a Context Block and a Feature Block, followed by a cross-modal fusion strategy to enhance prediction accuracy. Evaluated on a diverse dataset collected from the University of Leeds Driving Simulator, the proposed method achieves an accuracy of up to 95.8% in driver readiness classification. This performance significantly enhances existing approaches and highlights the importance of multimodal and multi-view fusion. As a real-time, non-intrusive solution, Driver-Net contributes meaningfully to the development of safer and more reliable automated vehicles and aligns with new regulatory mandates and upcoming safety standards.

## I. INTRODUCTION

Driver monitoring systems (DMS) are considered a critical part of new vehicles due to their role in enhancing the broader adoption of Autonomous Vehicles (AV) [1]. Such complex systems aim to monitor the driver’s level of readiness by detecting signs of distraction, fatigue, or driver impairments that can affect road safety.

The importance of DMS technology is constantly increasing, particularly as the safety regulatory organisations, such as the European New Car Assessment Programme (Euro NCAP), have defined a new set of mandatory tests and benchmarks for the inclusion of advanced driver assistance systems (ADAS), including DMS, to gain high safety ratings.

Since 2020, Euro NCAP has considered awarding points for vehicles equipped with systems capable of detecting driver inattention, distraction, or drowsiness. According to the European Union’s General Safety Regulation (GSR), every new vehicle manufactured from 7 July 2024 must be equipped with Driver Drowsiness and Attention Warning (DDAW) systems. By 7 July 2026, Advanced Driver Distraction Warning (ADDW) systems will also be mandatory [2]. The objective is not only to reduce accidents caused




Fig. 1. Proposed Cross-Modal fusion strategy for assessing driver situational awareness and readiness in response to critical take-over requests for Level-3 automated vehicles.

by human error, which accounts for over 90% of road crashes globally [3], but also to prepare the automotive industry for future roads dominated by semi-autonomous and autonomous vehicles.

The United States regulatory bodies, such as the National Highway Traffic Safety Administration (NHTSA), and Asian safety organisations try to align with such new safety requirements and standards which further emphasises the global importance of DMS technologies [4]. State-of-the-art research in DMS incorporates the latest developments in sensor technologies and utilises computer vision and machine learning techniques to achieve a robust and accurate driver state estimation. While advancements in autonomous vehicle technology have led to the deployment of robotaxis such as Waymo, in geofenced roads like San Francisco and Los Angeles, the primary focus of car manufacturers and academic researchers remains on enhancing the reliability of SAE Level 2 (L2) and Level 3 (L3) vehicles on public roads. According to the SAE automation levels [27], drivers in L2 vehicles with partial automation must continuously monitor the driving environment and retain full responsibility for vehicle control. In contrast, L3 vehicles, with conditional automation, allow drivers to disengage from active driving and engage in non-driving-related tasks (NDRTs) [28]. However, this modal shift can diminish the driver’s vigilance, lower their situational awareness, and pose potential challenges during critical take-over requests (TOR) [29], [30]. To address such challenges, DMS is tasked with continuously assessing the driver’s state in L3 vehicles to ensure safe and timely transitions from automated to manual control. This transition must safely occur within a few seconds, without worsening the complexity of the situation. If the DMS detects

<sup>†</sup> Corresponding author: Mahdi Rezaei (m.rezaei@leeds.ac.uk)

<sup>1</sup>M. Rezaei and M. Azarmi are with the Institute for Transport Studies, Computer Vision and Machine Learning Group, University of Leeds, LS2 9JT, United Kingdom (m.rezaei@leeds.ac.uk) and (tsmaz@leeds.ac.uk).

\* This work was funded and supported by the European Union’s Research and Innovation Programme and European Commission for the Hi-Drive Project.

that the driver is not adequately prepared for the take-over, the system will initiate a minimum risk manoeuvre [31], to stop the car in a safe location.

Among various driver state monitoring solutions, computer vision-based models are proven to be a promising solution in Level 3 vehicles to estimate driver readiness in response to take-over requests [19], particularly when the AV reaches the end of its operational design domain (ODD) or fails to operate in a complex scenario such as a road maintenance or lane blockage due to an accident. In the next section, we provide a literature review of the latest research in the field using in-cabin cameras and vision-based models.

## II. LITERATURE REVIEW

Modern DMS consider hybrid solutions, by using facial feature analysis [5], eye-tracking [6], and head pose estimation [7], to extract signs of driver distraction or fatigue. Convolutional Neural Networks (CNNs) [9] and Recurrent Neural Networks (RNNs) [8] have been extensively used for real-time video analysis to monitor yawning frequency or eyelid closures as key indicators of drowsiness. For example, Qu et al. developed a CNN-based spatio-temporal model to classify drivers' level of distraction, for a 10-class non-driving related task [20]. Similarly, Deo and Trivedi [22], along with Kazemi et al. [23], used Long Short-Term Memory (LSTM) networks to estimate driver take-over readiness, by utilising in-cabin visual cues and expert-labelled Observable Readiness Index (ORI).

To overcome the limitations of single-modality input, multimodal frameworks have been proposed. These systems extract various data types from wearable sensors, in-cabin cameras, and vehicle dynamics to provide a comprehensive perspective of the driver's state [14]. For example, Abbas et al. [21] combined multi-view cameras and physiological sensors within a CNN-RNN hybrid to detect driver inattention. Deng et al. [47] explored physiological responses such as heart rate (HR), skin conductance level (SCL), and mental workload (MWL), under simulated take-over conditions to quantify readiness delays. The "DeepTake" framework [26] extended this approach by predicting response time, intention, and take-over quality using a comprehensive set of vehicle and biometric features, although such wearable sensor-based systems face real-world scalability issues.

Some recent DMS research works have focused on near-infrared imaging and high-resolution cameras to detect micro-expressions and improve their robustness to lighting variations [10]. Deep learning models with attention mechanisms [11] and explainable architectures [12] are also being developed to improve the accuracy and interpretability of DMS. Furthermore, federated learning [17], [18] has emerged to address in-cabin privacy concerns while maintaining high model performance across diverse user datasets.

Several studies have focused on facial features and the driver's eye gaze. A recent study presented techniques for simultaneously detecting fatigue and distracted driving by utilising facial alignment networks to identify facial feature points and calculate distances to detect eye and mouth




Fig. 2. The University of Leeds Driving Simulator (top) with a Triple-camera setup for driver state monitoring (bottom) in a real Jaguar car inside the simulator dome.

states, achieving improved accuracy and computation time compared to previous approaches [24]. Bonyani et al. [32] propose the DIPNet model to predict driver intention that can be used as a sign of the driver's state in response to a TOR. Braunagel [45] explores an in-depth analysis of driver readiness in critical TOR via eye-on-the-road and gaze shift tracking in a simulator environment. Lui et. al. [48] propose the ACTNet model to predict minimum anticipated collision time (min ACT), as an indicator of drivers' readiness and performance for a take-over. The proposed CNN model mainly relies on eye-tracker data with a focus on eye fixation, blink rate, and pupil diameters as the main input features to determine driver readiness. However, reliance on eye-tracking data alone may not fully capture the multidimensional nature of readiness.

Other studies introduced large-scale datasets like the Driver Monitoring Dataset (DMD) [25], designed for benchmarking attention and drowsiness under various driving contexts. Transformer-based architectures [13], [15], [16] have also been applied to improve temporal feature modelling and cross-modal alignment.

In another study by Kim et al. [46], they propose a fundamental model for driver state estimation and investigate the influence of NDRTs on subjective driver readiness and take-over performance in L3 vehicles.

Despite these advancements, the reviewed works suffer from at least one of the following key limitations:

- **Mono-modality:** Many systems rely exclusively on facial or head-based inputs, and ignore critical cues from body posture and hand activity, which are vital for assessing true readiness.
- **Sensor Practicality:** High dependency on intrusive biosensors limits deployment feasibility in the AVs.
- **Spatio-Temporal Oversight:** Few studies model readi-



ness as a gradual, temporally evolving process. Our studies show that as the driver's attention reorients from non-driving tasks, the transition period typically spans over 1.6 seconds, to encompass steering wheel adjustments, situational scanning, and posture correction, which indicates the necessity of temporal modelling.

To address these gaps, we provide four key contributions as follows:

1- *Multi-Camera Cross-Modal Fusion Architecture*: We propose a first-of-its-kind triple-camera driver monitoring system that combines front, side, and over-shoulder views. The fusion of head pose, hand position, and body posture significantly outperforms single-modality approaches in assessing readiness.

2- *Spatio-Temporal Feature Integration*: Our method captures readiness as a dynamic, time-dependent process using synchronised video inputs and attention-based temporal modelling, moving beyond static indicators like gaze input.

3- *High-Performance, Real-Time Readiness Classification*: Driver-Net achieves a very high accuracy, demonstrating strong generalisation across diverse driver profiles and activities. The system is also lightweight enough for real-time deployment in modern vehicles.

4- *Safety-Critical Advancement*: By improving the accuracy and timing of take-over readiness detection, the proposed system supports safer transitions from automated to manual control, aligning with Euro NCAP and EU GSR 2024–2026 safety mandates and reducing risks associated with delayed driver responses.

### III. METHODOLOGY

The proposed Driver-Net deep neural network architecture processes  $N$ -frame temporal video sequences captured at 30 FPS from three synchronised HD cameras ( $1920 \times 1080$  pixels): front, side, and over-shoulder as depicted in Figure 3. The Raw video frames are then converted to  $720 \times 480$  pixels and are processed using three pre-trained feature extraction algorithms for hand detection, head-pose estimation, and body joint/pose estimation before being fed into the second layer of the model. The spatio-temporal video contents and extracted features are processed by two parallel 'Context' and 'Feature' blocks, and their outputs are fused by the 'Cross-Modal Fusion' block to finally predict the driver's readiness for a take-over request. Figure 3 illustrates the model inputs, feature extractors, and the network architecture, including the input-output sizes of each module.

#### A. Context Block

The Context Block is responsible for processing multi-view spatio-temporal features from input video sequences and aggregating them into a unified representation. A pre-trained ResNet-18 3D [33] prepares the input of the Context Block, which serves as a visual encoder to extract spatio-temporal features from all three input videos. The input tensor of the visual encoder has the shape  $\mathbb{R}^{B \times V \times N \times C \times W \times H}$ , where  $B$  is the batch size,  $V$  is the number of camera

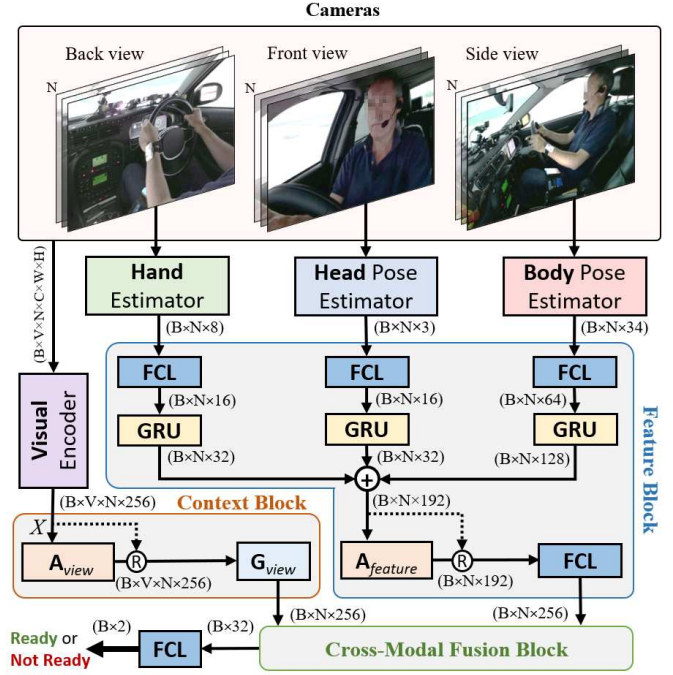


Fig. 3. The block diagram of the Driver-Net architecture, including camera input module and triple feature extractors, Visual encoder, Feature Block, Context Block, and Cross-Modal Fusion block.

views (front, side, over shoulder),  $N$  is the temporal sequence length,  $C$  is the number of colour channels (RGB), and  $W$ ,  $H$  are the spatial dimensions. Since ResNet-18 3D requires input tensors of shape  $\mathbb{R}^{B \times N \times C \times W \times H}$ , the camera views are flattened into the batch size dimension, resulting in input tensors with shape  $\mathbb{R}^{(B \cdot V) \times N \times C \times W \times H}$ . After processing by the encoder, the output feature maps have shape  $\mathbb{R}^{(B \cdot V) \times N \times 256}$ , following global average pooling across the spatial dimensions ( $W \times H$ ). These feature maps are then reshaped to  $\mathbb{R}^{B \times V \times N \times 256}$  for the subsequent layers.

A multi-head self-attention layer  $A_{\text{view}}$  operates on the dimension  $V$ , to capture inter-camera dependencies at each time step  $N$  while preserving the spatial and temporal structure. The self-attention layer with  $h = 4$  heads can be formulated as:

$$A_{\text{view}}(X) = (\text{Att}_1 \oplus \text{Att}_2 \oplus \dots \oplus \text{Att}_h) W_O \quad (1)$$

where  $\oplus$  denotes the concatenation operator, The input  $X$  which is processed by the visual encoder with shape  $\mathbb{R}^{B \times V \times N \times 256}$ , and  $W_O \in \mathbb{R}^{(h \cdot d_v) \times 256}$  is a learnable weight matrix that projects the concatenated attention heads back to the original feature dimension of 256. Here,  $d_v = 256/h$ , and the attention mechanism for  $h$ -th head,  $\text{Att}_h$ , is defined as follows:

$$\text{Att}_h(Q_h, K_h, V_h) = \text{softmax} \left( \frac{Q_h K_h^T}{\sqrt{d_k}} \right) V_h, \quad (2)$$

where the *query*, *key*, and *value* projections for the  $h$ -th head are given by:

$$Q_h, K_h, V_h = X W_h^Q, X W_h^K, X W_h^V \quad (3)$$

$W_h^Q \in \mathbb{R}^{256 \times d_k}$ ,  $W_h^K \in \mathbb{R}^{256 \times d_k}$ ,  $W_h^V \in \mathbb{R}^{256 \times d_v}$  are the learnable projection matrices for queries, keys, and values, respectively. The input  $X$  is reshaped from  $\mathbb{R}^{B \times V \times N \times 256}$  into  $\mathbb{R}^{(B \cdot N) \times V \times 256}$  to compute the attention over the  $V$  dimension. We also set  $d_k = d_v = 256/h$  to ensure consistency when concatenating the outputs of all heads.

To stabilise training and preserve gradient flow, a residual connection ‘R’ adds the input to the output of the attention layer (see dashed line arrow), followed by a layer normalisation [38]. The result is then passed to an aggregation module  $\mathbf{G}_{\text{view}}$  that summarises the multi-view camera information across the dimension  $V$ , producing outputs of shape  $\mathbb{R}^{B \times N \times 256}$ . The following three common aggregation methods were evaluated to find the most suitable for the task:

- 1) **Global Average Pooling (GAP):** Computes the average features across the entire dimension  $V$  at each time step in the input sequence with the length  $N$ , reducing feature complexity while preserving temporal dependencies.

$$\mathbf{G}_{\text{view}}^{\text{GAP}}(X_{\text{nr}}) = \frac{1}{V} \sum_{v=1}^V X_{\text{nr},v} \quad (4)$$

where  $X_{\text{nr}} \in \mathbb{R}^{B \cdot N \times 256}$  is the normalised input to the aggregation module after the residual connection and layer normalisation.

- 2) **Weighted Sum:** Applies a set of learnable weights  $w_v$  to each camera view’s features and sums them up.

$$\mathbf{G}_{\text{view}}^{\text{WS}}(X_{\text{nr}}) = \sum_{v=1}^V w_v \cdot X_{\text{nr},v} \quad (5)$$

where  $w_v \in \mathbb{R}^V$  are learned parameters which are normalised to satisfy  $\sum_{v=1}^V w_v = 1$ , ensuring interpretability and stability.

- 3) **1D Convolution:** A 1D convolution is applied across the dimension  $V$ , producing a summarised output over the  $V$  camera views.

$$\mathbf{G}_{\text{view}}^{\text{Conv1D}}(X_{\text{nr}}) = \text{Conv1D}(X_{\text{nr}}) \quad (6)$$

where Conv1D uses convolutional filters of shape  $(k, 256)$ , with the kernel size  $k$  with stride one. This method allows for learning spatial patterns across multiple camera inputs.

The Context Block’s final output is a tensor  $X_{\text{context}} \in \mathbb{R}^{B \times N \times 256}$  containing the processed spatio-temporal feature maps of the driver from three different camera views. The  $X_{\text{context}}$  is one of two inputs of the Cross-Modal Fusion block.

## B. Feature Block

The Feature Block processes the input videos and extracts driver-centric features (see Figure 4) using three computer vision algorithms as follows:

- **Head Pose Estimator:** The 3-D direction of driver head attention, including Yaw, Pitch, and Roll angles, is estimated using 3DDFA [34] from the frontal view camera. The collected features will be stored in the angles’ tensor  $X_{hd} \in \mathbb{R}^{B \times N \times 3}$ , to be used later in the Feature Block.
- **Body Pose Estimator:** Body joint features and their coordinates tensor  $X_{bp} \in \mathbb{R}^{B \times N \times 34}$  are extracted using YOLO7-Pose [35] from the side-view camera. Additionally, a YOLOv7 module detects objects within the region of the driver’s hand, represented by the tensor  $O \in \mathbb{R}^{B \times N \times M \times 2 \times 2}$ , where  $M$  is the number of detected objects per frame, and  $2 \times 2$  represents the dimension of the bounding box coordinates: top left corner  $(x_t, y_t)$  and bottom right corner  $(x_b, y_b)$ .
- **Hand Detector:** A YOLO7-based hand detection model [35] re-trained on the *EgoHands* dataset [36], and used to identify the position of the driver’s hands relative to the steering wheel from the rear-view (over shoulder) camera. The hand detector generates a tensor of bounding box coordinates for the left  $X_{hl} \in \mathbb{R}^{B \times N \times 4}$  and right  $X_{hr} \in \mathbb{R}^{B \times N \times 4}$  hand. These will be then concatenated to  $X_{hp} \in \mathbb{R}^{B \times N \times 8}$

A fully connected layer (FCL) independently transforms these tensors into higher-dimensional spaces while preserving temporal structure. For the input feature tensor  $X_f \in \mathbb{R}^{(B \cdot N) \times d_{\text{in}}}$ , which is flattened to the batch size ( $B$ ) and the sequence length ( $N$ ) dimension, the FCL uses a learnable weight matrix  $W \in \mathbb{R}^{(B \cdot N) \times d_{\text{out}}}$  and a vector  $b \in \mathbb{R}^{d_{\text{out}}}$ , producing the transformed output:

$$\text{FCL}(Y) = \text{ReLU}(X_f W + b) \quad (7)$$

where  $Y \in \mathbb{R}^{(B \cdot N) \times d_{\text{out}}}$  is the higher-dimensional representation of the input features, ReLU is a non-linear activation function [39], and  $d_{\text{in}}, d_{\text{out}}$  are the input and output feature dimensions, respectively. The  $d_{\text{in}} \rightarrow d_{\text{out}}$  transformation is  $34 \rightarrow 64$ ,  $3 \rightarrow 16$ , and  $8 \rightarrow 16$ , for the body pose, head pose; and hand status, respectively. (See Figure 3).

Once FCL units process the body, head, and hand feature vectors, a Gated Recurrent Unit (GRU) [37] captures temporal dependencies across the sequence of each tensor. Then, a concatenation operator combines the multiple GRU-processed tensors along with feature dimension, creating a new tensor of  $\mathbb{R}^{B \times N \times 192}$ .

A feature-wise self-attention layer  $\mathbf{A}_{\text{feature}}$  then operates on the feature dimension (i.e. 192), and dynamically identifies the most relevant features while de-prioritising redundant ones. Residual connections and layer normalisation ensure stable gradient propagation. An additional FCL processes the output tensor to align feature dimensions for next modules.

The Feature Block’s final output is a tensor  $X_{\text{feature}} \in \mathbb{R}^{B \times N \times 256}$  containing the processed temporal feature maps of the driver’s body pose, head angles, and hand position.

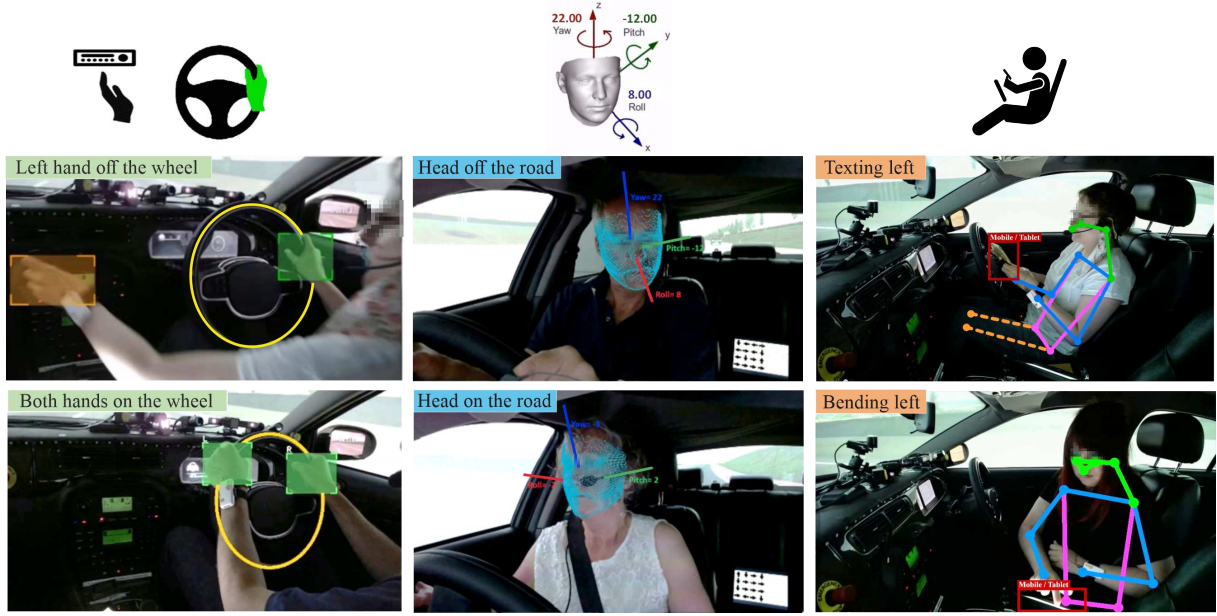


Fig. 4. **Left:** Hands/steering wheel positioning and activity recognition. **Middle:** Head-pose estimation (Yaw, Pitch, Roll  $\rightarrow$  head on/off the road). **Right:** Body pose estimation: normal, leaning/rotating, texting, and object detection.

### C. Cross-Modal Fusion Block

The Cross-Modal Fusion block integrates all spatio-temporal features from the Context and Feature blocks. Three widely used fusion methods in neural networks were evaluated:

- 1) **Concatenation Fusion (CF):**  $X_{\text{context}}$  and  $X_{\text{feature}}$  are concatenated along the feature dimension, yielding  $\mathbb{R}^{B \times N \times 512}$ . An FCL is then applied to process and reduce the input's feature dimensionality from  $\mathbb{R}^{B \times (N \cdot 512)}$  to  $\mathbb{R}^{B \times 256}$ . The FCL processes both the temporal and feature dimensions along the batches. CF is widely used in multimodal fusion for its simplicity and ability to preserve modality-specific features [41].
- 2) **Additive Fusion (AF):** This technique applies an element-wise summation of  $X_{\text{context}}$  and  $X_{\text{feature}}$  across the feature dimension, which results in a tensor  $\mathbb{R}^{B \times N \times 256}$ , followed by global average pooling across the temporal dimension  $N$ , producing a tensor with the shape  $\mathbb{R}^{B \times 256}$ . AF emphasises shared information and balances contributions of both modalities [42].
- 3) **Cross-Attention Fusion (CAF):** A cross-attention module, as defined in Eq. (1), computes the weights for both Feature and Context blocks, with  $X_{\text{context}}$  as the query and  $X_{\text{feature}}$  as the key and value. This enables the model to focus on the most relevant features and maximise the accuracy [43]. The outcome tensor is a fused representation of shape  $\mathbb{R}^{B \times 256}$ . CAF is particularly known for its effectiveness in fine-grained interaction modelling and has been widely applied in recent multimodal systems [44].

The output tensor of the fusion block is then processed and reduced to  $\mathbb{R}^{B \times 32}$  via an FCL with a 50% dropout rate to prevent overfitting [40].

Ultimately, the last FCL is used for the final classification to predict binary labels (“ready” or “not-ready”) for the driver state in response to a TOR request.

## IV. EXPERIMENTS

This section details the test environment, dataset, training and testing procedures, and the outcomes.

### A. Test Environment:

We conducted a series of experiments to assess the impact of each module variation on the overall performance of the proposed model. These experiments involved training and testing various components — namely, the Context Block, Feature Block, and Cross-Modal Fusion block — using the University of Leeds Driving Simulator (UoLDS), currently the most advanced driving simulator in the UK. The simulator dome includes a real Jaguar car and offers almost equal degrees of movement freedom with an immersive feeling of driving in the real world. The Jaguar cabin is equipped with a triple-camera module, positioned to capture comprehensive in-cabin driver behaviour, as illustrated in Figure 4:

- **Frontal Camera:** Positioned flexibly around the central dashboard, this camera faces the driver's head and feeds into the “Head Pose Estimator”, which infers the driver's visual attention using yaw, pitch, and roll angles, as shown in Fig. 4. It also tracks the duration of attention diversion, for example, by detecting sequences of  $N$  consecutive frames where the head orientation is off-road.
- **Over-Shoulder Camera:** Mounted to capture a rear view of the driver, this camera supports the “Hand Detector” module in identifying the position and activity of both hands. It determines whether one or both hands are on

the steering wheel, and if not, identifies their alternative locations and duration of disengagement—e.g., a hand on the gear lever or infotainment knob for  $t$  seconds before returning to the wheel.

- **Side Camera (Passenger-Side A-Pillar):** This camera captures the driver’s full-body posture and can identify contextual cues, such as interactions with handheld devices. The Body Pose Estimator uses this input to assess seating alignment (e.g., upright vs. leaning or reaching postures). Additionally, this side view provides redundancy in hand localisation, particularly valuable when occlusion affects the over-shoulder view, such as with tall drivers or those with larger body mass.

### B. Dataset:

We recruited 30 licensed drivers of various ages and sexes, all holding valid UK driving licences. The experiment was conducted in the University of Leeds Driving Simulator and encompassed both automated and manual driving modes on simulated multi-lane urban roads. During automated driving, participants were permitted to engage in NDRTs, such as using a tablet. In specific scenarios, participants received phone calls from the simulator control room and engaged in conversation while manually driving. Additionally, activities such as texting or browsing on a smartphone were intentionally introduced to simulate common distractions.

The UoLDS dataset consists of a selection of 540 video clips from each of the three in-cabin cameras, totalling 1620 videos. Clip durations range from 10 to 25 seconds, reflecting various driver behaviours across both automated and manual driving modes. The driver’s readiness status may change multiple times over time in the selected video cuts. However, only the last frame of each video clip was subjectively annotated by two expert reviewers as either “ready” or “not ready” in relation to a take-over request (TOR) initiated approximately two seconds earlier. These annotations were suggested by a thorough review of the synchronised triple-camera recordings, taking into account the driver’s overall in-cabin behaviour and situational context in the critical seconds preceding the TOR.

From the total of 1,620 annotated video samples, 654 instances (40.37%) were labelled as *ready*, and 966 instances (59.63%) were marked as *not ready*. These expert labels were subsequently used as ground truth for evaluating the model’s key performance indicators (KPIs). We used 1222 videos for training and 398 videos for testing. For the training phase, 50.16% of selected samples were labelled as ready, and 49.84% as not ready, demonstrating a balanced distribution of each label.

### C. Training and Test:

The model was trained and evaluated on a CUDA-enabled parallel computing platform equipped with an NVIDIA RTX A6000 GPU, 64GB of RAM, and an Intel Core i9-13900K 24-core processor.

The *Context* and *Feature* blocks were trained independently, each with a randomly initialised fully connected classification layer (FCL) to learn distinct representations. For the *Cross-Modal Fusion* block, we removed the classification heads from the pre-trained *Context* and *Feature* blocks, froze their learned parameters, and then only trained the *Cross-Modal Fusion* block with a randomly initialised classification head. Further implementation details are provided below.

TABLE I  
PERFORMANCE OF DIFFERENT AGGREGATION MODULES ( $G_{VIEW}$ ) FOR THE ‘CONTEXT BLOCK’ AGAINST ACCURACY (A), PRECISION (P), AND RECALL-RATE (R) METRICS.

#	Parameters	Aggregation	A(%)	P(%)	R(%)
$\alpha_{GAP}$	270,338	GAP	85.1	84.7	85.0
$\alpha_{WS}$	270,341	WS	84.5	84.2	84.4
$\alpha_{C1D}$	466,946	Conv1D	<b>88.0</b>	<b>87.4</b>	<b>87.7</b>

1) **Context Block Training:** The Context Block is aimed at learning a compact and discriminative representation of the driver’s behaviour from the last 1.6-second multi-view spatio-temporal features. The  $X_{context}$  tensor was flattened over the temporal dimension and connected directly to the classification head (final Fully Connected Layer, FCL) to predict driver readiness. Our experiments show that with a temporal dimension of  $N = 16$  (10 *fps* sampling), we can gain pretty much the same accuracy as the full-frame video input (30 *fps*). We evaluated three aggregation modules—Global Average Pooling (GAP), Weighted Sum (WS), and 1D Convolution (Conv1d) to aggregate the temporal features before classification. The Context block variants and their performance metrics on the test set are presented in Table I.

TABLE II  
PERFORMANCE OF EVERY SINGLE FEATURE AND ALL FEATURES COMBINED, THE ‘FEATURE BLOCK’ AGAINST ACCURACY (A), PRECISION (P), AND RECALL-RATE (R) METRICS.

#	Parameters	Feature set	A(%)	P(%)	R(%)
$\beta_P$	32,634	Body Pose	71.1	70.56	70.8
$\beta_A$	21,370	Head Angles	43.2	42.68	43.1
$\beta_H$	19,450	Hand Position	69.3	69.01	68.7
$\beta_{all}$	184,978	All Features	<b>86.94</b>	<b>86.30</b>	<b>86.50</b>

2) **Feature Block Training:** We trained the Feature Block aimed at the temporal learning of the driver’s state by integrating head pose, body pose, and hand position features extracted from a triple-camera setup. Similar to the Context Block, and in order to reduce the computation cost, the  $X_{feature}$  tensor was flattened over the dimension of  $N = 16$  and connected directly to the classification head to estimate the driver readiness state. We evaluated the performance of each feature set and the combined feature sets. The results are summarised in Table II, which proves the critical role of a triple-camera setup with a dramatic increase in the model performance in estimating driver readiness state.

3) **Cross-Modal Fusion Block:** Finally, we trained and evaluated the *Cross-Modal Fusion* block by combining the pre-trained *Context* and *Feature* Blocks. Three different



TABLE III  
PERFORMANCE OF DIFFERENT FUSION STRATEGIES FOR THE  
CROSS-MODAL FUSION BLOCK.

Model	Parameters	Fusion	A(%)	P(%)	R(%)	Cost
$\alpha_{GAP}, \beta_{all}$	2,536,592	CF	88.4	88.1	88.3	52ms
	447,152	AF	90.8	91.8	90.7	22ms
	1,036,976	CAF	92.8	93.1	91.8	35ms
$\alpha_{WS}, \beta_{all}$	2,536,595	CF	87.9	87.3	87.4	52ms
	447,155	AF	89.1	88.8	89.0	22ms
	1,036,979	CAF	90.8	91.8	90.8	35ms
$\alpha_{C1D}, \beta_{all}$	2,741,394	CF	<b>95.1</b>	<b>95.6</b>	<b>95.1</b>	62ms
	651,954	AF	91.2	91.9	<b>92.0</b>	24ms
	1,241,778	CAF	<b>93.2</b>	<b>93.7</b>	91.5	37ms

fusion strategies were trained: Concatenation Fusion (CF), Additive Fusion (AF), and Cross-Attention Fusion (CAF). The results are presented in Table III.

Table III shows that Concatenation Fusion and Cross-Attention Fusion perform best and second best, respectively, with the 1D convolution (in the Context block) and utilisation of all features (from the Feature block). Furthermore, the computational cost column ranging from 22ms to 62ms shows that the proposed model can process the input videos, ranging from 16 to 45 *fps*, which is fast enough for the intended real-world application of TOR.

To evaluate the generalisation capability of the proposed models, we performed 5-fold cross-validation ( $k = 5$ ) on the entire dataset (Table IV). Two training strategies were examined:  $\mathcal{T}_{all}$ , where the entire model including Context Block ( $\alpha_{C1D}$ ), Feature Block ( $\beta_{all}$ ), and Fusion Block were trained end-to-end and from scratch; and  $\mathcal{T}_{fusion}$ , where the Context block and Feature Block were frozen after pre-training, and only the Fusion Block was trained. As shown in Table IV,  $\mathcal{T}_{all}$  consistently outperforms  $\mathcal{T}_{fusion}$  in terms of accuracy, precision, and recall in line with the results obtained in Table III. Similarly, the Concatenation Fusion (CF) strategy performed best; however, Additive Fusion (AF) took the second-best place, despite a significantly smaller number of parameters. While these confirm the critical role of Fusion Block in the model's overall performance, it may indicate that AF's simplicity allows a more effective generalisation across multiple folds when the model is trained from scratch, while CF and CAF's complexity might lead to overfitting within individual folds of cross-validation.

TABLE IV  
PERFORMANCE EVALUATION USING 5-FOLD CROSS-VALIDATION ( $k=5$ ).

Train	Fusion	A $\pm \sigma$ (%)	P $\pm \sigma$ (%)	R $\pm \sigma$ (%)
$\mathcal{T}_{all}$	CF	<b>95.4</b> $\pm$ 0.3	<b>95.8</b> $\pm$ 0.4	<b>95.3</b> $\pm$ 0.3
$\mathcal{T}_{all}$	AF	<b>93.9</b> $\pm$ 0.4	<b>94.3</b> $\pm$ 0.5	<b>93.2</b> $\pm$ 0.4
$\mathcal{T}_{all}$	CAF	93.5 $\pm$ 0.5	94.0 $\pm$ 0.6	91.8 $\pm$ 0.5
$\mathcal{T}_{fusion}$	CF	93.1 $\pm$ 0.4	93.6 $\pm$ 0.5	92.9 $\pm$ 0.6
$\mathcal{T}_{fusion}$	AF	92.3 $\pm$ 0.4	92.7 $\pm$ 0.5	91.6 $\pm$ 0.4
$\mathcal{T}_{fusion}$	CAF	91.7 $\pm$ 0.6	92.3 $\pm$ 0.6	91.0 $\pm$ 0.5

## V. CONCLUSION AND FUTURE DIRECTIONS

This study presented a novel framework for estimating driver readiness in response to take-over requests (TOR) by leveraging multi-view spatio-temporal features extracted from a triple-camera setup and a quasi-naturalistic dataset

collected from the University of Leeds Driving Simulator (UoLDS). The proposed model integrated a Context Block, Feature Block, and Cross-Modal Fusion block to capture and fuse in-cabin behavioural cues, to enhance take-over readiness estimation.

Experimental results highlighted the effectiveness of each model component and the overall fused approach. The Context Block, which extracts multi-view contextual features, achieved the highest performance when using a 1D convolutional aggregation strategy, reaching an accuracy of 88.0%. The Feature Block, which utilises body pose, head angles, and hand position features, demonstrated significant improvements when combining all three modalities, achieving an accuracy of 86.94%. Finally, the Cross-Modal Fusion Block, responsible for merging contextual and feature representations, led to noticeable improvement and achieved the highest overall accuracy of 95.8% using a concatenation-based fusion strategy with Conv1D-driven context extraction. These findings suggest that the proposed approach offers a promising direction for driver monitoring systems, particularly in Level 3 conditionally automated vehicles, where accurate estimation of take-over readiness is critical for safety and performance.

While Driver-Net demonstrated strong performance and generalisation in estimating driver readiness using a multi-camera fusion strategy, the system is evaluated primarily within a high-fidelity simulator environment, which may not fully capture the variability and unpredictability of real-world driving conditions. Future work can focus on validating the model in real-world driving scenarios with a broader demographic pool. Incorporating non-intrusive physiological indicators and investigating continuous readiness scoring rather than binary classification may also enhance practical deployment and improve robustness and scalability.

## DECLARATIONS

The authors declare no conflicts of interest and no financial or personal relationships that could potentially influence the interpretation or presentation of the findings.

## ACKNOWLEDGMENT

This research has received funding from the European Union's Horizon 2020 research and innovation programme, for the Hi-Drive project under grant Agreement No 101006664. The article reflects only the author's view, and neither the European Commission nor CINEA is responsible for any use that may be made of the information this document contains.

## REFERENCES

- [1] Y. Qu, H. Hu, J. Liu, Z. Zhang, Y. Li and X. Ge, "Driver State Monitoring Technology for Conditionally Automated Vehicles: Review and Future Prospects," in IEEE Transactions on Instrumentation and Measurement, vol. 72, 1-20, 2023.
- [2] J. Lundahl, "Steering the Future: An Overview of Current and Upcoming Regulations in Automated Driving: Version 0.5," Diva Portal, 2024.
- [3] World Health Organization, "Global status report on road safety," WHO, 2018.



- [4] National Highway Traffic Safety Administration (NHTSA), "Advanced Impaired Driving Prevention Technology," NHTSA Research, nhtsa.gov, 2024.
- [5] D. Kim, H. Park, T. Kim, et al. "Real-time driver monitoring system with facial landmark-based eye closure detection and head pose recognition", *Scientific Reports* 13, 18264, 2023.
- [6] P. Pillai, B. Balasingam, Y.H. Kim, C. Lee, and F. BiondiPillai, "Eye-gaze metrics for cognitive load detection on a driving simulator," *IEEE/ASME Transactions on Mechatronics* 27.4, 2134-2141, 2022.
- [7] B. Ahn, D. G. Choi, L. Park, and I. S. Kweon, "Real-time head pose estimation using multi-task deep neural network," *Robotics and Autonomous Systems*, 103, 1-12, 2018.
- [8] M. Sun, M. Tsujikawa, Y. Onishi, X. Ma, A. Nishino, and S. Hashimoto, "A neural-network-based investigation of eye-related movements for accurate drowsiness estimation," *Intl. Conf. Engineering in Medicine and Biology Society*, pp. 5207-5210, 2018.
- [9] M. Ahmed, S. Masood, M. Ahmad and A. A. Abd El-Latif, "Intelligent Driver Drowsiness Detection for Traffic Safety Based on Multi CNN Deep Model and Facial Subsampling," in *IEEE Trans. on Intelligent Transportation Systems*, vol. 23, no. 10, 19743-19752, 2022.
- [10] E. Magdalena Nowara, T. K. Marks, H. Mansour, and A. Veeraraghavan, "SparsePPG: Towards driver monitoring using camera-based vital signs estimation in near-infrared," In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops* pp. 1272-1281, 2018.
- [11] F. Rahman and M. Chowdhury, "Enhancing DMS with explainable attention mechanisms," *Artificial Intelligence*, 2023.
- [12] M. G. Cimino, A. Di Tecco, P. F. Foglia, T. Nocchi, C. A. Prete, M. Ralli, and L. Tonelli, "Driver Distraction Detection on Edge Devices via Explainable Artificial Intelligence," *Multi-Conference on Systemics, Cybernetics and Informatics*, pp. 95-102, 2023.
- [13] Z. Hu, Y. Zhang, Y. Xing, Y. Zhao, D. Cao and C. Lv, "Toward Human-Centered Automated Driving: A Novel Spatiotemporal Vision Transformer-Enabled Head Tracker," in *IEEE Vehicular Technology Magazine*, vol. 17, no. 4, pp. 57-64, 2022.
- [14] A. Nair, V. Patil, R. Nair, A. Shetty, and M. Cherian. "A review on recent driver safety systems and its emerging solutions," *International Journal of Computers and Applications*, 46(3), pp. 137-151. 2024.
- [15] J. Yu, S. Park, S. Lee, M. Jeon, "Driver drowsiness detection using condition-adaptive representation learning framework," *IEEE transactions on intelligent transportation systems*, 20(11), 4206-4218. 2019.
- [16] Z. Li, L. Chen, L. Nie, and S. X. Yang, "A novel learning model of driver fatigue features representation for steering wheel angle". *IEEE Transactions on Vehicular Technology*, 71(1), 269-281. 2022.
- [17] R. Du, K. Han, R. Gupta, S. Chen, S. Labi and Z. Wang, "Driver Monitoring-Based Lane-Change Prediction: A Personalized Federated Learning Framework," 2023 *IEEE Intelligent Vehicles Symposium (IV)*, Anchorage, AK, USA, pp. 1-7, 2023.
- [18] Y. Cui, J. Zhu, and J. Li, "FLAV: Federated Learning for Autonomous Vehicle privacy protection," *Ad Hoc Networks*, 166, 103685, 2025.
- [19] V. Melcher, S. Rauh, F. Diederichs, H. Widloirther and W. Bauer, "Take-over requests for automated driving," *Procedia Manufacturing*, 3, 2867-2873, 2015.
- [20] F. Qu, N. Dang, B. Furht and M. Nojournian, "Comprehensive study of driver behavior monitoring systems using computer vision and machine learning techniques," *Journal of Big Data*, 11(1), 32, 2024.
- [21] Q. Abbas, M. E. Ibrahim, S. Khan and A. R. Baig, "Hypo-driver: a multiview driver fatigue and distraction level detection system," *CMC-computers Mater Contin*, 71(1), 1999-2017, 2022.
- [22] N. Deo and M. M. Trivedi, "Looking at the driver/rider in autonomous vehicles to predict take-over readiness," *IEEE Transactions on Intelligent Vehicles*, 5(1), 41-52, 2019.
- [23] M. Kazemi, M. Rezaei, M. Azarmi, "Evaluating Driver Readiness in Conditionally Automated Vehicles From Eye-Tracking Data and Head Pose," *IET Intelligent Transport Systems*, 19(1), 2025.
- [24] Y. Zheng, S. Chen, J. Wu, K. Chen, T. Wang and T. Peng, "Real-Time Driver Fatigue Detection Method Based on Comprehensive Facial Features," In *International Conference on Algorithms and Architectures for Parallel Processing*, 484-501, Springer Nature, 2023.
- [25] J.D. Ortega, N. Kose, P. Cafiias, M. A. Chao, A. Unnervik, M. Nieto and L. Salgado, "DMD: A large-scale multi-modal driver monitoring dataset for attention and alertness analysis," In *Computer Vision-ECCV 2020 Workshops: Glasgow, UK, August 23-28, 2020, Proceedings, Part IV* 16, pp. 387-405, 2020.
- [26] E. Pakdamanian, S. Sheng, S. Baee, S. Heo, S. Kraus, and L. Feng, "Deeptake: Prediction of driver takeover behavior using multimodal data," In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* pp. 1-14, 2021.
- [27] Society of Automotive Engineering, "SAE Levels of Driving Automation Refined for Clarity and International Audience.", Online: Retrieved from <https://www.sae.org/blog/sae-j3016-update>, January 2025.
- [28] X. Liu, R. Madigan, E. Sadraei, Y. M. Lee, and N. Merat, "Drivers' engagement in NDRTs during automated driving linked to travelling speed and surrounding traffic," *Transportation research part F: traffic psychology and behaviour*, 101, 332-339.
- [29] T. Louw, J. Kuo, R. Romano, V. Radhakrishnan, M. G. Lenné, and N. Merat, "Engaging in NDRTs affects drivers' responses and glance patterns after silent automation failures," *Transportation research part F: traffic psychology and behaviour*, 62, 870-882, 2019.
- [30] A. McKerral, K. Pammer, and C. Gauld, "Supervising the self-driving car: Situation awareness and fatigue during highly automated driving," *Accident Analysis & Prevention*, 187, 107068. 2023.
- [31] M. Cummings, "Assessing readiness of self-driving vehicles," In *The 103rd Transportation Research Board (TRB) Annual Meeting*, Washington DC, 2024.
- [32] M. Bonyani, M. Rahmian, S. Jahangard, and M. Rezaei, "DIPNet: Driver intention prediction for a safe takeover transition in autonomous vehicles," *IET Intelligent Transport Systems*, 17(9), 2023.
- [33] C. Feichtenhofer, H. Fan, J. Malik, and K. He, "Slowfast networks for video recognition," In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 6202-6211, 2019.
- [34] Z. Wang, X. Zhu, T. Zhang, B. Wang, and Z. Lei, "3D Face Reconstruction with the Geometric Guidance of Facial Part Segmentation," In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* pp. 1672-1682. 204.
- [35] C. Y. Wang, A. Bochkovskiy, and H. Y. M. Liao, "YOLOv7: Trainable bag-of-freebies sets new state-of-the-art for real-time object detectors," In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 7464-7475, 2023.
- [36] S. Bambach, S. Lee, D. J. Crandall, and C. Yu, "Lending a hand: Detecting hands and recognizing activities in complex egocentric interactions," *Intl. Conf. on Computer Vision*, pp. 1949-1957, 2015.
- [37] J. Chung, C. Gulcehre, K. Cho, and Y. Bengio, "Empirical evaluation of gated recurrent neural networks on sequence modeling," *arXiv preprint arXiv:1412.3555*, 2014.
- [38] J. Lei Ba, J. R. Kiros, and G. E. and G. E. Hinton, "Layer normalization", *ArXiv e-prints*, arXiv-1607, 2016.
- [39] V. Nair and G. E. Hinton, "Rectified linear units improve restricted Boltzmann machines," In *Proceedings of the 27th international conference on machine learning ICML-10*, pp. 807-814, 2010.
- [40] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever and R. Salakhutdinov, "Dropout: a simple way to prevent neural networks from overfitting," *Machine learning research*, 15(1), 1929-1958, 2014.
- [41] F. Zhao, C. Zhang, and B. Geng, "Deep Multimodal Data Fusion. *ACM Computing Surveys*, " 56(9), 1-36, 2024.
- [42] K. Bayoudh, R. Knani, R., F. Hamdaoui, A. Mtibaa, "A survey on deep multimodal learning for computer vision: advances, trends, applications, and datasets. *The Visual Computer*," 38(8), 2022.
- [43] J. Zhang, Y. Xie, W. Ding and Z. Wang, "Cross on cross attention: Deep fusion transformer for image captioning," *IEEE Transactions on Circuits and Systems for Video Technology*, 33(8), 4257-4268, 2023.
- [44] H. Akbari, L. Yuan, R. Qian, W. H. Chuang, S. F. Chang, Y. Cui, and B. Gong, "Vatt: Transformers for multimodal self-supervised learning from raw video, audio and text," *Advances in Neural Information Processing Systems*, 34, 24206-24221, 2021.
- [45] C. Braunagel, "Ensuring the take-over readiness of the driver based on the gaze behavior in conditionally automated driving scenarios," *Doctoral dissertation, Universität Tübingen*, 2018.
- [46] J. Kim, W. Kim, H. s. Kim, S. Lee, O. C. Kwon and D. Yoon, "A novel study on subjective driver readiness in terms of non-driving related tasks and take-over performance," *ICT Express*, 8(1), 91-96, 2022.
- [47] M. Deng, A. Gluck, Y. Zhao, D. Li, C. C. Menassa, V. R. Kamat, and J. Brinkley, "An analysis of physiological responses as indicators of driver takeover readiness in conditionally automated driving," *Accident Analysis & Prevention*, 195, 107372, 2024.
- [48] W. Liu, Q. Li, W. Wang, Z. Wang, C. Zeng and B. Cheng, "Deep Learning Based Take-Over Performance Prediction and Its Application on Intelligent Vehicles," in *IEEE Trans. on Intelligent Vehicles*, 2024.