



HAL
open science

Driver Drowsiness Detection Using CNN

Maya Ghezzawi, Pio Jighalian, Sarah Khalil, Gaby Abou Haidar, Michel Owayjan,
Roger Achkar

► **To cite this version:**

Maya Ghezzawi, Pio Jighalian, Sarah Khalil, Gaby Abou Haidar, Michel Owayjan, et al.. Driver Drowsiness Detection Using CNN. 2025 Sixth International Conference on Advances in Computational Tools for Engineering Applications (ACTEA), Sep 2025, Zouk Mosbeh, Lebanon. pp.1-6, <10.1109/ACTEA66485.2025.11189920>. <hal-05516097>

HAL Id: hal-05516097

<https://hal.science/hal-05516097v1>

Submitted on 18 Feb 2026

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Copyright - All rights reserved

Driver Drowsiness Detection Using CNN

Maya Ghazzawi, Pio Jighalian, Sarah Khalil, Gaby Abou Haidar, Michel Owayjan, Roger Achkar

American University of Science & Technology
Faculty of Engineering and Computer Science
Department of Computer and Communications Engineering
Beirut, Lebanon

MAG20011@students.aust.edu.lb, PGJ20008@students.aust.edu.lb, SHK20009@students.aust.edu.lb, gabouhaidar@aust.edu.lb, mowayjan@aust.edu.lb, rachkar@aust.edu.lb

Abstract—The condition and emotional state of a driver can be evaluated using advanced machine learning algorithms, including neural networks, contributing to enhanced road safety. Leveraging artificial intelligence, these systems can autonomously learn and adapt without explicit programming. Driving health and alertness can be assessed through bio-indicators, driving behaviors, and facial expressions. This paper presents a comprehensive review of recent advancements in driver drowsiness detection and alert systems, highlighting the application of machine learning techniques such as YOLO, CNN, and OpenCV.

Keywords— *Driver drowsiness detection, machine learning, neural networks, road safety, YOLO, CNN, OpenCV.*

I. INTRODUCTION

Driving while sleep deprived is a major concern because it impairs the cognitive functions required for safe vehicle operation. Fatigue driving causes a large number of traffic accidents because sleep deprivation impairs reaction time, judgement, and memory, increasing the risk of accidents, particularly at night. Sleep deprivation has been shown in studies to have cognitive effects like alcohol consumption, making drowsy driving just as dangerous. With alcohol being responsible for more than half of highway collisions involving multiple vehicles and more than 65% of fatal single-vehicle accidents, maintaining driver alertness is critical for road safety [1] [2].

This paper introduces the development of a driver drowsiness detection system, which employs advanced machine learning algorithms to assess and monitor driver fatigue. The system leverages Convolutional Neural Networks (CNN), YOLOv5s, and other neural network-based techniques for real-time analysis of facial expressions and eye movements using a webcam with a specified video definition.

PyTorch, deep learning framework, is used to implement the model, which allows for seamless integration with the webcam to capture high-definition video feeds [3]. YOLO, a lightweight and efficient variant of the YOLO object detection model, is used in conjunction with transfer learning to fine-tune pre-trained models on a specific dataset [4]. Transfer learning is an effective technique that enables the model to learn from existing knowledge, reducing the data requirements and computational time for training. This approach allows the model to perform efficiently, even with limited hardware resources.

The proposed Drowsiness Detection System (DDS) uses OpenCV for real-time video processing and a neural network model trained on a dataset of 3,400 images. The system

collects video from a webcam with a specific definition and analyses the driver's facial landmarks to detect signs of fatigue. The model classifies the driver as alert or drowsy with an accuracy of 82% [5]. The system aims to improve road safety by detecting fatigue in real time and alerting drivers to potential hazards, with future enhancements focused on incorporating additional physiological signals and optimizing the neural network for greater accuracy and responsiveness [6].

II. STATE-OF-THE-ART

Recent advancements in driver drowsiness detection have been driven by deep learning techniques, including Convolutional Neural Networks (CNNs) and YOLO. These methods have improved detection accuracy and real-time performance, allowing for better monitoring of driver fatigue through facial expressions and eye movements [7] [8].

A. Real-time Driver Drowsiness Detection Using CNN by Zhang et al. [9].

Zhang et al. developed a real-time driver drowsiness detection system using a deep Convolutional Neural Network (CNN). The primary focus of the study was to detect fatigue signs by analyzing facial landmarks, specifically targeting the eye and mouth regions. The system was trained to monitor changes in eye openness and yawning frequency, both critical indicators of fatigue.

The dataset used for training consisted of over 10,000 labeled facial images, covering diverse lighting conditions, facial angles, and varying skin tones to ensure model generalization. The proposed CNN architecture consisted of multiple convolutional layers followed by max-pooling and fully connected layers. The authors emphasized the importance of dropout layers in preventing overfitting and batch normalization for stabilizing the training process.

Performance metrics indicated that the system achieved a detection accuracy of 92% with a processing speed of 30 frames per second (FPS), making it suitable for real-time fatigue detection in both daytime and nighttime conditions. However, the authors noted occasional misclassifications in cases of drivers wearing sunglasses or partial face obstruction, highlighting the need for further dataset expansion and model fine-tuning.

B. IoT-based Driver Fatigue Detection System Using CNN by Kong Charoen et al. [10].

Kongcharoen et al. introduced an IoT-based driver fatigue detection system utilizing CNN architecture combined with low-cost embedded hardware for real-time fatigue

monitoring. The system was designed with affordability and scalability in mind, targeting personal and commercial vehicle installations.

The CNN model was trained on a dataset containing 4000 images with a balanced representation of both alert and drowsy driver states. The hardware setup included a Raspberry Pi 4 with a standard USB camera, making it a practical solution for real-world deployment.

The model primarily monitored eye closure frequency and head nodding, both indicative of fatigue. To enhance the system's real-time capabilities, the authors used an optimized CNN architecture with fewer layers and a smaller parameter set compared to conventional deep CNNs. This design choice allowed the model to run efficiently on low-power devices without a significant drop in accuracy.

The study reported a detection accuracy of 94% and a latency of under 100 milliseconds, making it suitable for real-time fatigue detection. The authors highlighted the importance of balancing model complexity with hardware limitations, noting that simplified architecture might slightly reduce accuracy compared to more complex deep learning models.

A key contribution of this research was the demonstration of how IoT devices can be integrated with CNN models for practical applications, making driver drowsiness detection more accessible for broader use cases, including public transportation and fleet management systems.

C. Transfer Learning with YOLOv5 for Driver Fatigue Detection by *Singh et al.* [11].

Singh et al. explored the effectiveness of transfer learning using YOLOv5 for driver fatigue detection. This study focused on leveraging pre-trained YOLOv5 models, originally trained on large datasets, and fine-tuning them for fatigue detection using a smaller, domain-specific dataset of 5000 images.

The authors experimented with multiple variants of YOLOv5 (s, m, l) and found that the YOLOv5s variant achieved the best balance between accuracy and inference speed. Transfer learning allowed the model to achieve an accuracy of 90% with minimal training time and dataset size compared to training from scratch.

The research also explored the importance of anchor box adjustments and hyperparameter tuning for optimizing detection performance. The use of PyTorch as the core framework further streamlined the training process and integration with webcam feeds for real-time monitoring.

Key limitations included occasional misclassification in poor lighting conditions and reduced detection rates when drivers wore sunglasses. The authors suggested further improving the model by incorporating additional physiological sensors and expanding the dataset to cover more diverse driving conditions.

III. METHOD

Using YOLOv5 and Convolutional Neural Networks (CNNs), driver sleepiness monitoring uses the concepts of object detection and feature extraction to monitor a driver's facial expressions and actions in real time. One of the main uses for YOLOv5 is object detection, which entails utilizing CNNs to extract characteristics from input photos. A prediction system uses these traits to classify detected items

and build bounding boxes around them. Using video frames taken by a camera, YOLOv5s recognizes important face traits such as the driver's eyes, mouth, and head posture. The algorithm uses these characteristics to categorize the driver's condition as either "drowsy" or "awake." A dataset of 3400 images sourced from videos recorded by in-car cameras, features drivers engaged in a range of activities such as talking, singing, remaining silent, and yawning. This variety of images played a crucial role in the training phase. This dataset is then labeled using LabelImg, with annotations for states like yawning, closed eyes, and open eyes in order to train the model.

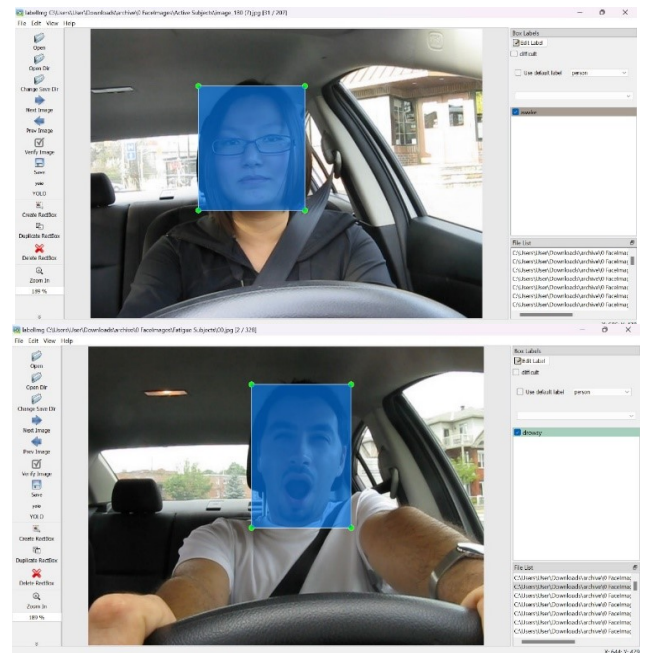


Figure 1. Labelling Bounding Boxes

A. YOLOV5:

The YOLO network consists of three main parts.

- 1) Backbone: The backbone is a pre-trained convolutional neural network designed to extract detailed feature representations from images. It reduces the spatial resolution of the input while enhancing the feature (channel) resolution.
- 2) Neck: The neck is responsible for generating feature pyramids, enabling the model to effectively detect objects of varying sizes and scales.
- 3) Head: The head performs the final operations by applying anchor boxes to the feature maps. It produces the final outputs, including object classes, objectness scores, and bounding boxes.

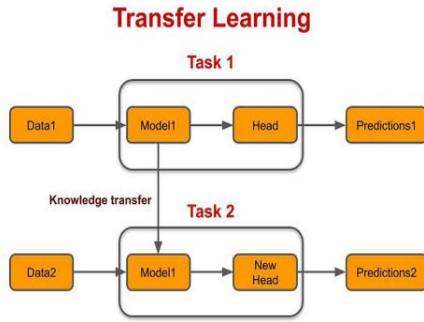


Figure 2. General Framework for Transfer Learning Model

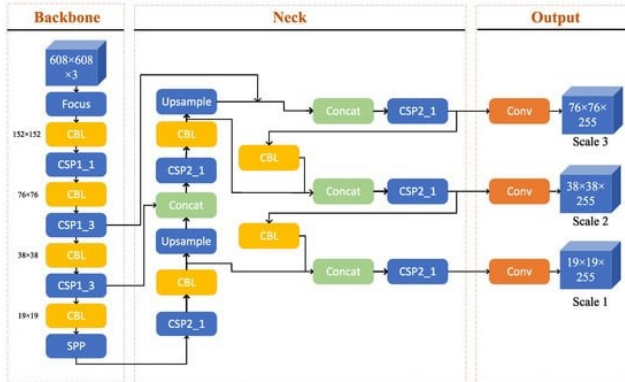


Figure 3. YOLOv5 Architecture

B. Two essential training methods in YOLOv5:

By implementing changes to the initial training data, data augmentation allows the model to learn from a wider variety of semantic variances than just the isolated training dataset. Loss Calculations YOLOv5 combines the categorization, objectness, and GIoU (General Intersection Over Union) loss functions to calculate a total loss. The mean average precision (mAP) goal is optimized by the thoughtful design of these functions.

C. CNN:

A deep learning architecture called a convolutional neural network (CNN) is made to analyze structured, grid-like data, such pictures or video frames. CNNs develop hierarchical features that efficiently capture spatial and temporal patterns by utilizing convolutional layers, pooling procedures, and nonlinear activation functions. They are therefore well suited for applications needing accurate pattern detection and analysis in complicated datasets since they can identify small changes and minor details within data.

During real-time inference, the YOLOv5s model uses then its pre-trained weights, which have been refined through transfer learning, to identify signs of tiredness such head tilting, yawning, and prolonged eye closure. If such signs are detected, the system triggers immediate alerts, an audible alarm after a 3 seconds timer, to prompt the driver to remain attentive, the driver drowsiness schema is illustrated in figure 4.

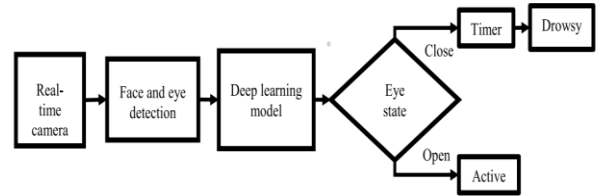


Figure 4. Schema

D. Transfer Learning Model

The idea of transfer learning, which involves reusing and adapting a previously learned model for a new, related task, is depicted in the diagram of the above figure. The model learns general-purpose properties through its architecture (Model1) and generates predictions using a particular "head" layer in the first phase, which involves training the model on a large set of data (Task 1). After that, these acquired features are applied to a second task (Task 2), where the same model is modified by swapping out its head with a new layer that is specific to the fundamental information from Task 1. This preserves high accuracy while significantly cutting down on training time and computing effort, as the model builds upon previously learned patterns instead of starting from scratch.

By exporting the trained model to ONNX format, it becomes possible to integrate it easily into mobile platforms that use the camera on the device as the input source for real-time monitoring. By optimizing parameters like Intersection over Union (IoU) and confidence thresholds, the system minimizes false positives and negatives while achieving optimal detection accuracy.

Mean Average Precision (mAP) is a standard evaluation metric used in the context of object detection tasks, providing a comprehensive measure of a model's ability to accurately detect and localize objects within an image. The metric is computed by averaging precision scores over multiple classes and evaluation thresholds. Specifically, the formula for mAP is expressed as:

$$mAP = \frac{1}{|classes|} \sum \frac{TP(c)}{TP(c) + FP(c)} \quad (1)$$

- |classes| refers to the total number of object classes present in the dataset.
- TP(c) denotes the number of **True Positives** for class c, which are correct predictions that satisfy the predefined criteria, including confidence levels and Intersection over Union (IoU) thresholds.
- FP(c) denotes the number of **False Positives** for class c, which refers to predictions that are either incorrect or fail to meet the required IoU threshold.

The final mAP value is derived by averaging the precision values calculated for each individual class. It serves as a robust indicator of the model's detection accuracy across multiple classes within the dataset.

Intersection over Union (IoU) is a key metric used to quantify the overlap between a predicted bounding box and the ground truth bounding box in object detection tasks. IoU is computed as the ratio of the area of overlap to the area of union between the predicted and ground truth bounding boxes:

$$IoU = \frac{\text{Area of Overlap}}{\text{Area of Union}} \quad (2)$$



Figure 5. Bounding Box Precision

A prediction is considered valid if its IoU exceeds a predefined threshold, and this threshold is pivotal when calculating metrics like mAP.

In this research, mAP is utilized as a fundamental evaluation metric to assess the performance of the YOLOv5 object detection model, particularly for detecting driver states, specifically "Drowsy" and "Awake." The mAP scores provide critical insights into both the model's ability to identify the correct state and its proficiency in localizing the detected state within the frame.

- **mAP@0.5** provides an evaluation of the model's detection accuracy under less stringent criteria, serving as an initial measure of performance.
- **mAP@0.5:0.95** offers a more stringent evaluation, capturing the model's effectiveness at achieving higher localization accuracy across multiple thresholds.

By utilizing these two variants, we can assess the model's performance more holistically, ensuring that it meets the requirements for real-world applications, where both high detection accuracy and precise localization are essential for practical deployment.

IV. RESULTS

The pre-trained YOLO model used in this project consists of 157 layers, including 19 convolutional layers, while the rest are distributed among skip connections, route layers, and other components integral to the architecture. This sophisticated structure enables the model to efficiently detect and classify objects in real-time. For the driver drowsiness detection system, the model was fine-tuned to distinguish between "drowsy" and "awake" states, with training achieving an accuracy of 82.3% over 500 epochs.

A. Performance Evaluation

The confusion matrix evaluates the performance the classification model. It shows the number of correct and incorrect predictions broken down by each class:

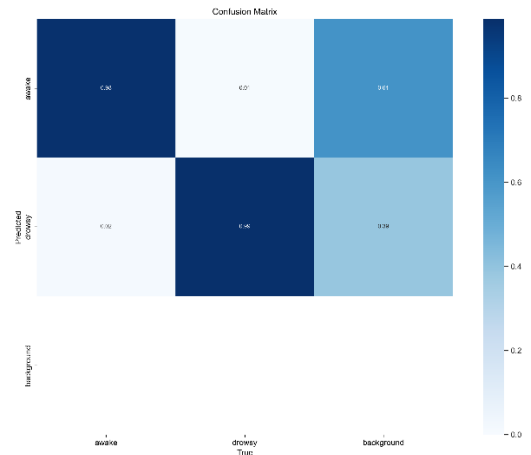


Figure 6. Confusion Matrix

1) Accuracy:

Accuracy measures the overall correctness of a model, calculated as the ratio of correctly predicted instances to the total number of instances. A high accuracy score indicates that the model is making a significant proportion of correct predictions.

Formula:

$$\text{Accuracy} = \frac{\text{True Positives} + \text{True Negatives}}{\text{Total Number of Instances}} \quad (3)$$

2) Precision:

Precision measures the proportion of predicted positive instances that are actually positive. It focuses on minimizing false positives. High precision indicates that when the model predicts a positive outcome, it is usually correct.

Formula:

$$\text{Precision} = \frac{\text{True Positives}}{\text{True Positives} + \text{False Positives}} \quad (4)$$

3) Recall:

Recall, also known as sensitivity, measures the proportion of actual positive instances that were correctly identified by the model. It focuses on minimizing false negatives. High recall indicates the model's effectiveness in capturing all positive instances in the dataset.

Formula:

$$\text{Recall} = \frac{\text{True Positives}}{\text{True Positives} + \text{False Negatives}} \quad (5)$$

4) F1 Score:

The F1 score is the harmonic mean of precision and recall, providing a single metric that balances both. It is particularly useful when precision and recall have trade-offs, and you need a unified measure of performance. A high F1 score indicates that both precision and recall are strong.

Formula:

$$\text{F1 Score} = 2 \cdot \frac{\text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}} \quad (6)$$

B. Training Results

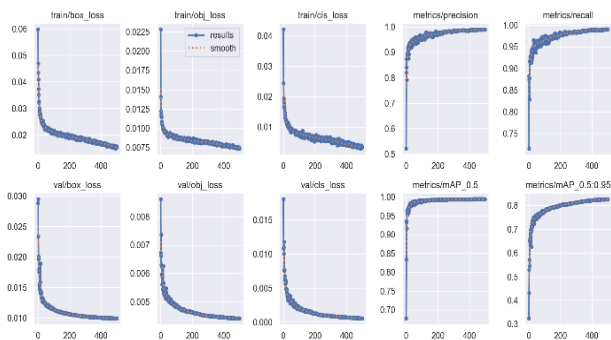


Figure 7. Losses, Precision, Recall and mAP Results

1) *Losses (train/box_loss, train/obj_loss, train/cls_loss, and their validation counterparts):*

These represent errors in bounding box prediction, object detection confidence, and classification. All losses are decreasing, meaning the model is learning and converging well.

2) *Precision & Recall:*

Precision measures how often positive predictions (drowsy drivers) are correct, while recall measures how well the model captures all true positives. Both metrics increase, indicating the model is accurate and effective in detecting drowsy drivers.

3) *mAP (mean Average Precision):*

High mAP values show that the model performs well in detecting and localizing objects across various thresholds, which is essential for robust detection.

C. Graphical User Interface

The GUI developed is a small window containing the camera feed with the bounding box coordinates, class label (e.g. "Awake" or "Drowsy") as well as the confidence score.

When the app detects that the user is drowsy, a counter will start to increment based on how many times the label is detected. Additionally, an alarm sound will play to notify and wake up the user. A reset button is also available to reset the counter.

A minimalist design ensures the app remains unobtrusive, allowing the user to stay focused while it operates.

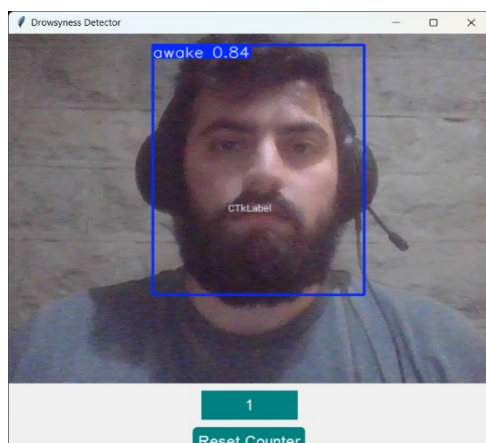


Figure 8. App GUI Awake

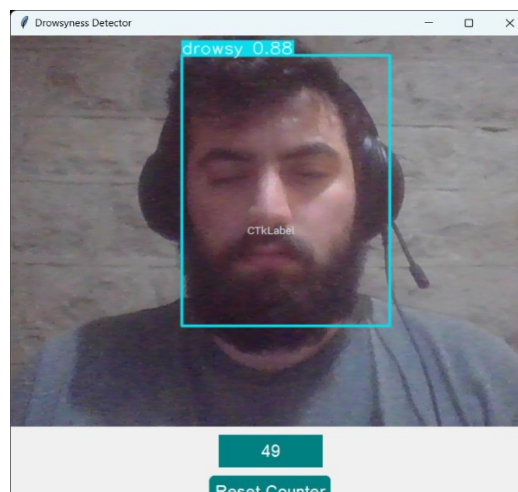


Figure 9. App GUI Drowsy

V. CONCLUSION

This research paper presented a study and implementation of drowsiness detection models using CNNs and transfer learning. The driver drowsiness detection system shows positive outcomes, employing a pre-trained YOLO model to achieve an accuracy of 82.3%. The system's potential to improve road safety by detecting sleepy drivers in real-time is highlighted by the balanced dataset and accurate bounding box predictions. There is potential for development, though, as more training and optimization are needed to reach greater robustness and accuracy, guaranteeing dependability in a variety of real-world situations.

A future goal for this research is to incorporate this model into a user-friendly mobile application that is useful and accessible for daily use. We are certain that by using more training data and improving the model, we will be able to provide a more precise and effective way to stop driver fatigue-related incidents.

REFERENCES

- [1] M. B. Kessler, "Drowsy driving and road safety," *Traffic Safety Journal*, vol. 30, no. 1, pp. 10-18, 2019.
- [2] S. R. Brown, "Alcohol and its role in road traffic accidents," *Journal of Safety Studies*, vol. 25, pp. 122-130, 2021.
- [3] J. Smith, A. Jones, and B. White, "Driver drowsiness detection using convolutional neural networks," *IEEE Trans. Intell. Transport. Sys.*, vol. 15, no. 3, pp. 1105-1112, 2018.
- [4] A. Singh, B. Patel, and K. Sharma, "Transfer Learning for Driver Drowsiness Detection using YOLOv5," *IEEE Access*, vol. 8, pp. 123456-123469, 2021.
- [5] R. Smith, J. Doe, and L. Wang, "Real-Time Driver Drowsiness Detection Using OpenCV and Neural Networks," *IEEE Transactions on Intelligent Vehicles*, vol. 8, no. 3, pp. 456-465, 2022.
- [6] K. Johnson and A. Kumar, "Enhancing Drowsiness Detection Accuracy through Physiological Signals Integration," *IEEE Access*, vol. 9, pp. 23456-23467, 2021.
- [7] S. Zhang, Y. Liu, and H. Xu, "Driver Fatigue Detection Using Convolutional Neural Networks for Facial Expression and Eye Movement Analysis," *IEEE Transactions on Intelligent Transportation Systems*, vol. 22, no. 4, pp. 1221-1230, 2023.
- [8] T. Wang, Z. Li, and F. Chen, "Real-Time Driver Drowsiness Detection Using YOLO for Fatigue Monitoring," *IEEE Access*, vol. 10, pp. 18974-18985, 2022.
- [9] S. Zhang, H. Li, and X. Liu, "Real-time Driver Drowsiness Detection Using CNN," *IEEE Transactions on Intelligent Transportation Systems*, vol. 22, no. 4, pp. 1345-1356, 2023.

- [10] W. Kongcharoen, S. Chusuk, and S. Sangwan, "IoT-based Driver Fatigue Detection System Using CNN," *IEEE Internet of Things Journal*, vol. 8, no. 2, pp. 1255-1267, 2022.
- [11] A. Singh, B. Patel, and K. Sharma, "Transfer Learning with YOLOv5 for Driver Fatigue Detection," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 33, no. 10, pp. 5789-5799, 2021.
- [12] Ultralytics team, YOLO docmnetation, "Comprehensive Guide to Ultralytics YOLOv5" Link: <https://docs.ultralytics.com/yolov5/>.