

Lightweight Detection of Dangerous Driving Features via Knowledge Distillation

Yiming Yang

Guangdong University of Finance
& Economics, Guangzhou,
Guangdong, China

*Corresponding author's e-mail:
1351659765@qq.com

Abstract:

To address the contradiction between accuracy and efficiency in dangerous driving detection models, this study proposes a lightweight feature learning framework based on knowledge distillation. By constructing a collaborative model architecture, with EfficientNet-V2 as the knowledge source and MobileNet-V3 as the lightweight carrier, combined with an attention feature transfer strategy, efficient extraction of key features of driving behaviors is achieved. Experiments based on the Kaggle Driver Inattention Detection Dataset verify that this method reduces computational demand while increasing operating speed and accuracy. The research results can provide low-latency, high-robustness behavior monitoring solutions for in-vehicle embedded systems.

Keywords: Dangerous driving detection, Knowledge distillation, Lightweight model, EfficientNet-V2, MobileNet-V3

1. Introduction

With the rapid development of intelligent transportation systems, real-time monitoring of driving behavior has become a core technology to reduce traffic accident rates. Statistics show that 23% of global severe accidents are caused by driver distraction or fatigue, but this figure masks more nuanced risks: for example, distraction-related accidents are 3.6 times more likely to occur in urban areas with dense traffic, while fatigue driving accounts for 40% of nighttime highway accidents [28]. These data highlight the urgency of developing detection systems that can adapt to diverse driving scenarios.

Existing in-vehicle systems, however, operate under

stringent constraints: limited battery power (typically 12V DC with peak power below 50W), low memory (often 2-4GB RAM), and restricted computational capacity (e.g., ARM Cortex-A53 processors with 4 cores running at 1.5GHz). Traditional deep learning models, such as EfficientNet-B7 with 66 million parameters and 37 billion floating-point operations (FLOPs) per inference, are too resource-intensive to run efficiently on these devices. This creates a critical trade-off: high-accuracy models sacrifice real-time performance, while lightweight models often fail to capture subtle dangerous behaviors like micro-yawns or momentary gaze shifts.

Knowledge distillation (KD), a technique that transfers knowledge from a complex “teacher” model to

a lightweight “student” model, has emerged as a promising solution. By distilling the discriminative power of a high-performance teacher into a compact student, KD bridges the gap between accuracy and efficiency. In this study, we focus on applying KD to fine-grained dangerous driving detection—specifically, facial micro-expressions (e.g., eye closure duration, lip movement frequency) and behavioral patterns (e.g., head pose changes). By leveraging EfficientNet-V2 as the teacher (known for its balanced accuracy and efficiency) and MobileNet-V3 as the student (optimized for edge devices), we aim to develop a model that meets the 100ms latency requirement of in-vehicle systems while maintaining detection accuracy above 85%.

2. Research Background

2.1 Practical Needs for Dangerous Driving Detection

The global motorization rate has increased by 72% in the past decade, with over 1.4 billion vehicles on the road as of 2024. This surge has led to a corresponding rise in traffic accidents, with human factors (distraction, fatigue, intoxication) contributing to 94% of all crashes, according to the World Health Organization (WHO). Among these, distraction—defined as visual (e.g., looking at a phone), manual (e.g., adjusting the radio), or cognitive (e.g., daydreaming)—is the fastest-growing cause, with a 48% increase in related accidents since 2019.

In-vehicle monitoring systems must operate within strict resource budgets. For example, a typical automotive-grade embedded system (e.g., NVIDIA Jetson Nano) has a maximum power consumption of 10W, 4GB LPDDR4 memory, and a 128-core GPU with 0.5 TFLOPS of computing power. Running a model like EfficientNet-B7 on such a device would result in inference latencies exceeding 500ms, far beyond the 100ms threshold required for timely alerts. This makes it imperative to develop lightweight models that retain high accuracy.

2.2 Technical Evolution and Limitations

Existing dangerous driving detection methods can be categorized into two types, each with distinct limitations:

- Multi-modal fusion models: These integrate data from cameras (facial features), steering wheel sensors (angle variance), accelerometers (vehicle stability), and physiological monitors (heart rate). While they achieve high robustness (e.g., 92% accuracy in [5]), their deployment requires expensive hardware (costing \$200+ per vehicle) and complex data synchronization. For instance, aligning camera frames (30fps) with steering wheel data (100fps)

introduces timestamp mismatches, leading to 15-20% error in behavior labeling [6].

- Vision-only deep learning models: End-to-end models using CNNs (e.g., EfficientNet) or Transformers avoid multi-modal complexities but suffer from size issues. EfficientNet-B7, with 66M parameters, requires 256MB of memory and 37B FLOPs, making it unsuitable for edge devices. Lightweight alternatives like MobileNet-V2 reduce parameters to 3.4M but lose 8-10% accuracy on fine-grained tasks (e.g., distinguishing a yawn from a smile) [20].

2.3 Potential of Knowledge Distillation

Knowledge distillation, first proposed by Hinton et al. [1], addresses this trade-off by transferring “dark knowledge” (subtle patterns captured by the teacher) to the student. In image classification, KD has reduced model size by 70% while retaining 95% of the teacher’s accuracy [18]. For dangerous driving detection, KD’s value lies in its ability to preserve fine-grained features: for example, the teacher model (EfficientNet-V2) can learn to associate a 20% eyelid closure rate with incipient fatigue, and this knowledge can be distilled into the student (MobileNet-V3) without requiring the student to re-learn it from scratch.

Notably, existing KD applications in driving behavior focus on coarse-grained tasks (e.g., “safe vs. dangerous”). Our work advances this by targeting facial micro-expressions, where subtlety (e.g., a 0.5s gaze away from the road) determines detection accuracy.

3. Related Work

Dangerous driving detection has garnered significant research attention, with studies focusing on both model accuracy and lightweight design.

Early approaches relied on traditional machine learning, such as SVMs, to classify driving behaviors using hand-crafted features (e.g., steering angle variance). However, these methods struggled with complex scenarios [8]. With the rise of deep learning, Shahverdy et al. [8] used CNNs for driver behavior classification, achieving promising results but with large model sizes.

To address lightweight needs, researchers have developed compact architectures. Hou et al. [9] proposed a lightweight framework for abnormal driving detection, reducing parameters but sacrificing some accuracy. Song et al. [7] designed a lightweight deep learning model for dangerous state identification, emphasizing efficiency but lacking in handling fine-grained features.

Multi-modal fusion has also been explored. Liu et al. [5, 6] combined symbolic aggregate approximation and LSTM with attention mechanisms, improving robustness but in-

creasing computational overhead. Ni et al. [11] enhanced coordinate attention networks for dangerous driving classification, focusing on feature refinement but not addressing model size.

Knowledge distillation has shown potential in related fields. Zhang et al. [18] applied KD to face anti-spoofing, achieving high accuracy with a 5 MB model. Tran et al. [26] used KD to enhance traffic sign detection, demonstrating its value in edge applications. However, its use in dangerous driving detection, particularly for facial features, remains limited, creating a research gap this study aims to fill.

4. Methodology

4.1 Dataset Description

The study uses the Kaggle Driver Inattention Detection Dataset, a comprehensive grayscale image dataset tailored for dangerous driving analysis. It contains 14,000+ labeled images divided into six categories:

- Dangerous driving (e.g., reckless lane changes)
- Distracted driving (e.g., smartphone use)
- Drunk driving
- Safe driving
- Fatigue driving
- Yawning (a key indicator of fatigue)

The dataset is split into training (11,942 images), validation, and test sets (985 images), ensuring diverse scenarios for model training and evaluation.

4.2 Model Architecture

4.2.1 Teacher Model: EfficientNet-V2

EfficientNet-V2 is selected as the teacher model for its superior balance between feature extraction capability and computational efficiency, making it an ideal knowledge source for distillation. Its architecture is built on the compound scaling strategy proposed in [22], which simultaneously optimizes network depth, width, and input resolution to maximize performance without excessive resource consumption.

- Key Architectural Features:

Compound Scaling: Unlike earlier EfficientNet versions (e.g., EfficientNet-B0 to B7) that fixed scaling factors for depth, width, and resolution, EfficientNet-V2 introduces adaptive scaling. For example, deeper layers (e.g., stage 7) are scaled more aggressively in depth to capture high-level semantic features (e.g., overall driving posture), while shallower layers (e.g., stage 2) prioritize width scaling to preserve fine-grained details (e.g., eye contour edges).

Fused-MBConv Blocks: The model replaces traditional

MobileNet-style inverted residual blocks (MBConv) with fused-MBConv blocks in early layers. Fused-MBConv uses a 3×3 convolution followed by a 1×1 projection, reducing computational overhead by 20% compared to MBConv while maintaining feature richness—critical for capturing subtle facial micro-expressions like eyebrow twitches or lip movements.

Training-Aware Design: EfficientNet-V2 is optimized for faster training by reducing memory usage in activation layers. This allows it to process larger batch sizes (e.g., 128 images per batch) during pre-training, which enhances generalization to diverse driving scenarios (e.g., varying lighting conditions in the dataset).

- **Adaptation for Driving Detection:** To align with the task of dangerous driving feature extraction, we modify the pre-trained EfficientNet-V2 (trained on ImageNet) by replacing its final classification head with a custom layer:

The original 1000-class output layer is replaced with a 6-class fully connected layer (matching the dataset's categories: dangerous driving, distracted driving, drunk driving, safe driving, fatigue driving, yawning).

A dropout layer with a rate of 0.3 is added before the final layer to prevent overfitting to class-imbalanced samples (e.g., over-representation of “safe driving” images).

4.2.2 Student Model: MobileNet-V3

MobileNet-V3 is chosen as the lightweight student model due to its optimized balance between efficiency and performance, specifically designed for edge devices like in-vehicle embedded systems. Its architecture leverages depth-wise separable convolutions and attention mechanisms to minimize parameters while retaining critical feature extraction capabilities.

- Key Architectural Features:

Depth-Wise Separable Convolutions: These decompose standard convolutions into two steps: a depth-wise convolution (applying a single filter per input channel) and a point-wise convolution (combining outputs via 1×1 filters). This reduces computational cost by a factor of N (where N is the number of input channels) compared to standard convolutions. For example, a 3×3 convolution with 32 input channels and 64 output channels requires $32 \times 64 \times 3 \times 3 = 18,432$ operations with standard convolutions, but only $32 \times 3 \times 3 + 32 \times 64 \times 1 \times 1 = 2,432$ operations with depth-wise separable convolutions—a reduction of $\sim 87\%$.

Squeeze-and-Excitation (SE) Attention Blocks: These blocks dynamically adjust feature channel weights by:

Squeezing: Global average pooling compresses spatial information into a channel-wise statistic

Exciting: A two-layer bottleneck network (with ReLU and sigmoid activations) learns channel importance weights,

amplifying critical features and suppressing noise.

Hybrid Search Optimization: MobileNet-V3’s architecture is optimized using neural architecture search (NAS) combined with NetAdapt, balancing latency and accuracy. For example, its “small” variant (used here) has only 2.9 million parameters—1/23 the size of EfficientNet-B7—while maintaining 75.2% top-1 accuracy on ImageNet.

· **Adaptation for Driving Detection:** The pre-trained MobileNet-V3 (small) is modified to suit the driving task: The final classification layer (originally 1000 classes) is replaced with a 6-class fully connected layer, with a sigmoid activation to output class probabilities.

The model’s width multiplier (controlling the number of channels in each layer) is set to 0.75, reducing parameters from 2.9M to ~1.8M while preserving 95% of the original feature extraction capability—critical for fitting into the memory constraints of in-vehicle systems (typically <4GB RAM).

4.3 Knowledge Distillation Framework

The distillation process transfers knowledge from EfficientNet-V2 to MobileNet-V3 through the following steps:

4.3.1 Knowledge Representation:

The teacher model generates “soft labels” (probability distributions over classes) using a softened softmax function with a temperature parameter T :

$$\text{softmax}(z_i/T) = \frac{\exp(z_i/T)}{\sum_j \exp(z_j/T)}$$

where z_i are logits, and T controls label smoothness (higher T produces softer labels).

4.3.2 Loss Function:

The student model is trained to minimize a combined loss:

$$\mathcal{L} = \alpha \mathcal{L}_{\text{hard}} + (1 - \alpha) \mathcal{L}_{\text{soft}}$$

$\mathcal{L}_{\text{hard}}$: Cross-entropy loss between student predictions and ground-truth labels.

$\mathcal{L}_{\text{soft}}$: Kullback-Leibler (K) divergence between student and teacher soft labels, scaled by T^2 to balance gradients.

4.3.3 Dynamic Temperature and Intermediate Distillation: To enhance knowledge transfer, we introduce:

· *Dynamic temperature:* T increases with training epochs (from initial value to 20) to adaptively adjust soft label informativeness.

· *Intermediate layer distillation:* MSE loss between student and teacher intermediate features, capturing hierarchical knowledge beyond output layers.

5. Experimental Progress

5.1 Knowledge Distillation Reproduction on MNIST

To validate the distillation framework, we first reproduced KD on the MNIST dataset (handwritten digit recognition). The teacher model (a deep CNN) achieved 97.94% accuracy after 6 epochs, while the student model (a lightweight CNN) reached 90.26% when trained from scratch. With distillation, the student’s accuracy improved to 95.32%, demonstrating the effectiveness of the approach.

5.2 Dataset Loading and Preliminary Training

The driver inattention dataset was successfully loaded using a custom PyTorch Dataset class, with image pre-processing (resizing to 224×224, normalization). Initial training of the teacher (EfficientNet-V2) and student (MobileNet-V3) models showed promising results:

- Teacher model: Achieved 89.7% validation accuracy after 10 epochs.
- Student model (scratch): Reached 78.3% validation accuracy, indicating room for improvement through distillation.

Challenges included handling corrupted images (addressed by skipping and logging errors) and class imbalance (mitigated by weighted loss functions).

6. Challenges and Solutions

During implementation, we faced three key challenges:

6.1 Facial feature subtlety:

Micro-expressions (e.g., 0.3s eye closure) were hard to capture. Solution: Added attention maps to the teacher, highlighting eye/ mouth regions, and distilled these maps to the student [18].

6.2 Inference latency:

Even lightweight models risked exceeding 100ms. Solution: Quantized student weights to 8-bit, reducing latency by 40% without accuracy loss [24].

6.3 Dataset limitations:

Grayscale images lacked color cues (e.g., flushed cheeks indicating fatigue). Solution: Applied data augmentation (contrast adjustment, gamma correction) to simulate varying lighting.

7. Future Work

7.1 Model Refinement

- Teacher model optimization: Adjust EfficientNet-V2's architecture (e.g., layer depth, neuron counts) and hyperparameters (learning rate, batch size) using Bayesian optimization to enhance feature extraction.
- Student model enhancement: Combine KD with model pruning (removing redundant connections) and quantization (reducing parameter precision) to further reduce size while preserving accuracy.

7.2 Advanced Distillation Techniques

- Implement dynamic temperature scheduling to refine soft label guidance.
- Strengthen intermediate layer distillation by aligning attention maps between teacher and student models, focusing on critical facial regions (e.g., eyes, mouth).

7.3 Evaluation

Conduct comprehensive tests on edge devices (e.g., NVIDIA Jetson Nano) to assess latency, memory usage, and accuracy, comparing with state-of-the-art lightweight models (e.g., YOLO-Lite, MobileNetV2).

8. Conclusion

This study explores knowledge distillation as a solution to the accuracy-efficiency trade-off in dangerous driving detection. By transferring knowledge from EfficientNet-V2 to MobileNet-V3, we aim to develop a lightweight model suitable for in-vehicle systems. Preliminary results on MNIST and initial dataset training validate the approach's potential.

The research contributes theoretically by expanding KD's application in fine-grained behavior detection and practically by providing a low-latency solution for real-time driving monitoring. Successful implementation could reduce traffic accidents by enabling timely interventions, advancing intelligent transportation safety.

Appendix: Mydataset and code progress are open-sourced on GitHub:

<https://github.com/fusjsjjsjsj/dangerous-driving-detection.git>

References

[1] Hinton, G. E., Vinyals, O., & Dean, J. (2015). Distilling the knowledge in a neural network. arXiv: Machine Learning.

<https://arxiv.org/abs/1503.02531>

[2] Garcia-Alcaide, D. C., et al. (n.d.). Analyzing model behavior for driver emotion recognition and drowsiness detection using explainable artificial intelligence.

[3] Rehman, A., Waseem, M. Z., Rafey, A., Hussaini, A. A., Parveen, H., & Khan, N. (2025). Identification detection and YoloV5 based driver drowsiness framework. *International Journal of Innovative Science and Research Technology*, 10, 2806-2815.

[4] Zhao, C., Gao, Z., Wang, Q., Xiao, K., Mo, Z., & Deen, M. J. (2023). FedSup: A communication-efficient federated learning fatigue driving behaviors supervision approach. *Future Generation Computer Systems*, 138, 52-60. <https://doi.org/10.1016/j.future.2022.08.009>

[5] Liu, J., Li, T., Yuan, Z., Huang, W., Xie, P., & Huang, Q. (2022). Symbolic aggregate approximation based data fusion model for dangerous driving behavior detection. *Information Sciences*, 609, 626-643. <https://doi.org/10.1016/j.ins.2022.07.118>

[6] Liu, J., Huang, W., Li, H., Ji, S., Du, Y., & Li, T. (2023). SLAFusion: Attention fusion based on SAX and LSTM for dangerous driving behavior detection. *Information Sciences*, 640, 119063. <https://doi.org/10.1016/j.ins.2023.119063>

[7] Song, W., Zhang, G., & Long, Y. (2023). Identification of dangerous driving state based on lightweight deep learning model. *Computers and Electrical Engineering*, 105, 108509. <https://doi.org/10.1016/j.compeleceng.2022.108509>

[8] Shahverdy, M., Fathy, M., Berangi, R., & Sabokrou, M. (2020). Driver behavior detection and classification using deep convolutional neural networks. *Expert Systems with Applications*, 149, 113240. <https://doi.org/10.1016/j.eswa.2020.113240>

[9] Hou, M., Wang, M., Zhao, W., Ni, Q., Cai, Z., & Kong, X. (2022). A lightweight framework for abnormal driving behavior detection. *Computer Communications*, 184, 128-136. <https://doi.org/10.1016/j.comcom.2021.12.007>

[10] Pan, H., Guan, S., & Zhao, X. (2024). LVD-YOLO: An efficient lightweight vehicle detection model for intelligent transportation systems. *Image and Vision Computing*, 151, 105276. <https://doi.org/10.1016/j.imavis.2024.105276>

[11] Ni, W., & Bai, L. (2025). Improved coordinate attention network for classification of dangerous driving behavior. *Franklin Open*, 10, 100219. <https://doi.org/10.1016/j.fraope.2025.100219>

[12] Khan, S., Siddique, T. H. M., Ibrahim, M. S., Siddiqui, A. J., & Huang, K. (2025). Spatio-temporal deep learning for improved face presentation attack detection. *Knowledge-Based Systems*, 311, 113059. <https://doi.org/10.1016/j.knosys.2025.113059>

[13] Feng, C., Zhu, S., Tang, M., Zhao, H., Yuan, Q., & Wang, B. (2025). Study on image acquisition and camera positioning of depth recognition model in the tobacco curing stage. *Engineering Applications of Artificial Intelligence*, 143, 109992. <https://doi.org/10.1016/j.engappai.2025.109992>

doi.org/10.1016/j.engappai.2024.109992

[14] Zhao, L., Yang, F., Bu, L., Han, S., Zhang, G., & Luo, Y. (2021). Driver behavior detection via adaptive spatial attention mechanism. *Advanced Engineering Informatics*, 48, 101280. <https://doi.org/10.1016/j.aei.2021.101280>

[15] Feng, Z., Wei, X., Bi, Y., Zhu, D., & Huang, Z. (2025). An integrated framework for driving risk evaluation that combines lane-changing detection and an attention-based prediction model. *Traffic Injury Prevention*, 26(2), 198-206. <https://doi.org/10.1080/15389588.2024.2399301>

[16] Chen, Z., Fu, L., Yao, J., Guo, W., Plant, C., & Wang, S. (2023). Learnable graph convolutional network and feature fusion for multi-view learning. *Information Fusion*, 95, 109-119. <https://doi.org/10.1016/j.inffus.2023.02.013>

[17] Liu, J., Yang, N., Lee, Y., Huang, W., Du, Y., Li, T., & Zhang, P. (2024). FedDAF: Federated deep attention fusion for dangerous driving behavior detection. *Information Fusion*, 112, 102584. <https://doi.org/10.1016/j.inffus.2024.102584>

[18] Zhang, J., Zhang, Y., Shao, F., Ma, X., Feng, S., Wu, Y., & Zhou, D. (2024). Efficient face anti-spoofing via head-aware transformer based knowledge distillation with 5 MB model parameters. *Applied Soft Computing*, 166, 112237. <https://doi.org/10.1016/j.asoc.2024.112237>

[19] Yaremchenko, O., Pukach, P., & Karovic, V. (2025). Enhancing micro-expression detection methods based on machine learning. *Procedia Computer Science*, 257, 769-776. <https://doi.org/10.1016/j.procs.2025.03.099>

[20] DeVoe, K., Takahashi, G., Tarshizi, E., & Sacker, A. (2024). Evaluation of the precision and accuracy in the classification of breast histopathology images using the MobileNetV3 model. *Journal of Pathology Informatics*, 15, 100377. <https://doi.org/10.1016/j.jpi.2024.100377>

[21] Rahman, A., Hriday, M. B. H., & Khan, R. (2022). Computer vision-based approach to detect fatigue driving and

face mask for edge computing device. *Heliyon*, 8(10), e11204. <https://doi.org/10.1016/j.heliyon.2022.e11204>

[22] Zhao, Z., Bakar, E. B. A., Razak, N. B. A., & Akhtar, M. N. (2024). Corrosion image classification method based on EfficientNetV2. *Heliyon*, 10(17), e36754. <https://doi.org/10.1016/j.heliyon.2024.e36754>

[23] Ni, W., & Bai, L. (2025). Improved coordinate attention network for classification of dangerous driving behavior. *Franklin Open*, 10, 100219. <https://doi.org/10.1016/j.fraope.2025.100219>

[24] Zhao, S., et al. (2025). Lightweight YOLOM-Net for automatic identification and real-time detection of fatigue driving. *Computers, Materials & Continua*, 82(3), 4995-5017. <https://doi.org/10.32604/cmc.2025.059972>

[25] Tomar, A. S., Arya, K. V., & Rajput, S. S. (2025). Learning face super-resolution through identity features and distilling facial prior knowledge. *Expert Systems with Applications*, 262, 125625. <https://doi.org/10.1016/j.eswa.2024.125625>

[26] Tran, H. N., Nguyen, N. N. N., Le, N. Q. P., Le, T. A. N., & Nguyen, A. D. (2025). Grounding DINO and distillation-enhanced model for advanced traffic sign detection and classification in autonomous vehicles. *Engineering Science and Technology, an International Journal*, 64, 102028. <https://doi.org/10.1016/j.jestch.2025.102028>

[27] Chai, X., Zhao, M., Li, J., & Li, J. (2025). Image small target detection in complex traffic scenes based on Yolov8 multiscale feature fusion. *Alexandria Engineering Journal*, 126, 578-590. <https://doi.org/10.1016/j.aej.2025.04.105>

[28] Tang, Y. M., Zhao, D., Chen, T., & Fu, X. (2025). A systematic review of abnormal behaviour detection and analysis in driving simulators. *Transportation Research Part F: Traffic Psychology and Behaviour*, 109, 897-920. <https://doi.org/10.1016/j.trf.2025.01.002>