

16TH ITU ACADEMIC CONFERENCE


ITUKALEIDOSCOPE

GENEVA2026

*AI and frontier
technologies for good*


7-9 July 2026
Geneva, Switzerland

Hosted by  **AI for Good**
Global Summit



Driver Drowsiness Detection on the Edge

Comparative evaluation of lightweight deep learning
models for on-device driver monitoring



ITU Kaleidoscope 2026 · Geneva,
Switzerland · 7-9 July 2026





**Zenon Lamprou, Georgia Christodoulou,
Konstantinos Avgerinakis**

Catalink Limited · ALFIE Project (Horizon Europe,
Grant No. 101177912)

ITU Kaleidoscope 2026 · Geneva, 7-9 July 2026



ITU KALEIDOSCOPE
GENEVA2026

Why It Matters: Safety & Regulation

DROWSY DRIVING: THE SCALE

91,000

US crashes annually (NHTSA)

10-20%

EU road accidents caused by fatigue

EU MANDATE: GSR 2019/2144


Jul 2024 – DDAMS Mandatory

All new EU vehicles must comply

2030 – -50% Deaths Target

2050 – Vision Zero

DDAMS = Driver Drowsiness and
Attention Monitoring System



Data & Models

DATASETS (~57K IMAGES)

Custom Dataset

16,788 imgs • Alert / Drowsy / Yawning

DDD - Driver Drowsiness Dataset

~41,000 imgs • 22+19 drivers

Yawn Detection Dataset

1,448 imgs • Binary yawning classification

Diverse race, gender & recording conditions

MODELS EVALUATED (6)

DozeGuardRLDD

SigLIP fine-tune • vision-language encoder

MobileViT

CNN × ViT hybrid • ~6M params


YOLO Classifier

YOLOv8-cls • Ultralytics framework

mini-Xception ★

Depth-sep. CNN • <1 MB • 58K params

★ *Top performer in cross-dataset evaluation*



Methodology & Key Results

EVALUATION PIPELINE

1. Collect & Sample

Balanced random sampling across datasets

2. Two Pre-processing Conditions

Cond. A: Full frame (no crop)

Cond. B: Face crop + CLAHE

3. Inference & Score

Accuracy + F1 on 3 datasets

4. LiteRT Export

On-device inference evaluation

BASELINE ACCURACY (NO CROP)

Model	DDD	Yawn	Custom
YOLO Cls	1.00	0.55	0.69
mini-Xception	0.68	0.52	0.91
MobileViT	1.00	0.65	0.73
DozeGuard	0.90	0.60	0.58

FACE CROP IMPACT (YAWN DET.)

MobileViT	0.65 → 0.96	(+0.31)
mini-Xcp	0.52 → 0.86	(+0.34)

DDD perfect scores → cross-dataset collapse



Edge Deployment & Conclusions

LITERT (TFLITE) CONVERSION

Why on-device?

Real-time latency • GDPR compliance
Privacy-preserving • Offline-capable

Conversion Result

mini-Xception & YOLO → ~0.50 acc (chance)

Root Causes

Operator incompatibilities in LiteRT
INT8/FP16 quantization precision loss
No unified conversion framework

MAIN CONCLUSIONS

Training data diversity is the most critical factor for real-world DMS deployment

Face crop + CLAHE consistently improves accuracy at zero overhead

mini-Xception: best cross-dataset generaliser (<1 MB)


Current LiteRT pipelines are inadequate for safety-critical edge detection

FUTURE WORK

Unified quantization-aware training for reliable edge conversion

Spatio-temporal models (3D CNN, ConvLSTM)

Fairness evaluation across demographic groups (ongoing)



Thank you!

Zenon Lamprou • Catalink Limited

ALFIE Project – Horizon Europe • Grant No. 101177912

<https://alfie-project.eu/>



ALFIE has received funding by the European Union's
Horizon Europe research and innovation programme
under Grant Agreement N° 101177912

