Deployment-Oriented Evaluation of YOLOv8 for Real-Time Face Detection in Unconstrained Video Streams

Prihartini Widiyanti*, Nabil Abdulkarim Ishaq, Ghulam Muhammad, Kencono Wungu, and Olivia Trisha Nethania[1]

[1]Master Study Program of Biomedical Engineering, Faculty of Science and Technology, Universitas Airlangga, Surabaya 60115, Indonesia
Corresponding author, E-mail: pwidiyanti@fst.unair.ac.id

Abstract

Real-time face detection is a fundamental task in modern computer vision, with applications in surveillance, security, healthcare monitoring, and human–computer interaction. While specialized face detectors achieve high benchmark accuracy, their computational complexity often limits their deployment on resource-constrained systems. This study evaluates the practical feasibility of YOLOv8 for real-time face detection and analyzes its performance under unconstrained real-world conditions. A YOLOv8-based face detection system was implemented and tested on continuous video streams captured in natural indoor environments using consumer-grade hardware. The evaluation focused on empirical detection performance and per-frame processing latency. The system achieved an average inference latency of approximately 48 ms per frame, corresponding to near real-time performance. High-confidence detections were consistently observed for frontal and near-frontal faces under moderate illumination, while detection reliability decreased under low-light conditions, partial occlusion, motion blur, and when faces occupied a small portion of the frame. These findings indicate that YOLOv8 provides a favorable balance between computational efficiency and practical detection reliability in deployment-oriented real-time scenarios, although robustness under challenging environmental conditions remains an area for further improvement.

Keywords: *YOLOv8, face detection, real-time video processing, deployment-oriented evaluation, object detection, deep learning, computer vision*