OPEN ACCESS
International Journal of Innovations in Science & Technology
RESEARCH & INNOVATION DIVISION
IJIST

No Smoke: Smart Surveillance for Cigarette Smoking Detection and Student Identification System

Zartasha Baloch, Irfan Ali Bhacho, Madeha Memon, Mahwish Soomro, Isra Samoo, Muzaffar Hussain

Department of Computer Systems Engineering, Mehran University of Engineering and Technology, Jamshoro Pakistan.

***Correspondence**: zartasha.baloch@faculty.muet.edu.pk

Citation | Baloch. Z, Bhacho. I. A, Memon. M, Soomro. M, Samoo. I, Hussain. M, "No Smoke: Smart Surveillance for Cigarette Smoking Detection and Student Identification System", IJIST, Vol. 8 Issue.3 pp1429-1448, June2026

DOI: https://doi.org/10.33411/IJIST/1936

Received | May 10, 2026 **Revised |** June12, 2026 **Accepted |** June16, 2026 **Published |** June25, 2026.

Cigarette smoking in educational institutes poses significant health and disciplinary challenges for university administration despite strict institutional policies. Manual surveillance systems are often inefficient in detecting such violations, especially within large campuses and indoor areas. To address this issue, this research presents an AI-based cigarette smoking detection and student identification system that leverages deep learning and computer vision for real-time surveillance. The proposed system uses YOLOv8 for cigarette detection from CCTV feeds. A custom dataset of cigarette-smoking images was used to train and validate the model to ensure robustness under diverse lighting and background conditions. The system then identifies the student using MTCNN and FaceNet for face detection. The system also integrates a web dashboard for an alert mechanism that captures screenshots when a cigarette-smoking violation is detected and notifies the concerned authorities via email to facilitate immediate action. The YOLOv8s cigarette detection module achieved an accuracy of 90.06%, precision of 95.80%, recall of 91.27%, F1-score of 93.48%, specificity of 85.71%, false positive rate of 14.29%, and a mAP@0.5 of 88%, demonstrating strong and reliable detection performance under diverse real-world CCTV conditions. The experimental results of the face recognition module have also shown satisfactory performance over all evaluation metrics, achieving an overall accuracy of 92.2%, precision of 91.8%, recall of 92.2%, and F1-score of 91.9%. The system contributes to maintaining smoke-free educational environments by enabling proactive policy enforcement. Overall, this research highlights the potential of AI-powered surveillance to reduce smoking incidents and foster a healthier academic community.

Keywords: Cigarette smoking; Computer Vision; Face Detection; Object Detection; Web Dashboard

Introduction:

Globally, the use of tobacco and nicotine products is rapidly increasing, particularly among the younger generation. There are many reasons for the widespread use of cigarettes that includes social environment and pressure, curiosity, lack of knowledge of its adverse impacts and ease of cigarette access [1]. Frequent smoking in teenagers causes health issues such as frequent respiratory infections, poor lung development, reduced lung capacity, and even lung cancer[2][3]. Alhindal et al. reviewed the adverse effects of smoking on cardiovascular system and addressed this problem in Adults as Congenital Heart Disease (ACHD) [4]. Currently university campuses lack an effective automatic system which detects cigarette smoking within their premises. As a result, anti-smoking policies are often violated frequently in areas such as corridors, libraries, classrooms, and other indoor spaces where smoking is strictly prohibited. These activities not only violate the university rules and regulations, but also increase health risks for not only smokers but also for non-smokers. This research focuses on the development of an Artificial Intelligence (AI) based surveillance system that uses cameras and intelligent detection algorithms to identify and report smoking activities in real-time. This system enforces anti-smoking policies more effectively and reduces unhealthy practices and promotes a healthier anti-smoking academic environment. The objectives of this research are:

To collect and preprocess datasets for cigarette and face detection.

To train and evaluate the model for cigarette detection using YOLOv8.

To enable accurate matching of detected faces against **the** enrolled student database using MTCNN and FaceNet models.

To design a web application that integrates the trained model.

Related Work:

The rise of AI-based surveillance in educational settings has drawn increasing attention to implement automated solutions for policy enforcement and safety monitoring. In a recent review by [5], the authors highlighted the addictive nature of nicotine from smoking or electronic vaping and advocated for AI-based surveillance in educational institutions. [6] found that 20.6% of the students were current smokers, with the majority being in high school (54.9%) and in university (23.5%). They identified higher smoking levels among male students aged 21 years or older in research sample. [7] explored the use of Machine Learning (ML) techniques to detect smoking behavior and provides evidence that AI-powered vision systems can reliably identify smoking attempts. The study supports implementing an AI-based surveillance system for indoor university spaces. [8] reviewed 49 studies published between 2019 and 2023 from PubMed and ACM Digital Library, focusing on technology features, outcomes, and evaluation methods. Machine learning models and wireless signal detection techniques yield encouraging results for smoking surveillance but require further refinement. The identification of these limitations along with continuous research and innovation in this field are crucial for developing reliable, practical solutions and integrating these technologies into clinical programs. [9] proposed a smoking detection model using the Single Shot Detector (SSD) method with VGG16 as the feature extractor. Experimental results demonstrated that at a threshold of 0.5, the model achieved outstanding performance, with 97.5% accuracy, 97.4% precision, 100% recall, and a 98.7% F1-score. [10]proposed the Context-Aware Small-Item Attention Net (CASIA-Net), a novel detection framework to address the issues caused due to the small scale of targets, poor visibility, and cluttered environments. CASIA-Net incorporates a Deformable Feature Extraction (DFE) module to tackle the non-salient characteristics of smoking targets.

[11] developed a face-authentication system by merging the Deep Learning (DL) models like MTCNN (Multi-Task Cascade Convolutional Network) for accurate face detection and for high precision facial extraction and face embeddings using FaceNet. This

paper provides the base for real time monitoring systems, it also discusses the key challenges such as variations in pose, lighting, and occlusion that must be kept in mind while designing AI-based smoking detection systems for campus use.

[12] developed a real-time cigarette detection using YOLOv3 object detection algorithm that is trained on a dataset of 10.000 images where (5,000 cigarette images vs 5,000 non-cigarette images) under different variations of lighting conditions, angles and backgrounds. [13] evaluated the performance of different YOLO variants on the smoker detection problem using small objects such as cigarettes and introduced a new dataset (CigDet) published in the Mendeley Data repository, highlighting challenges around dataset availability and size for accurate deep learning model training.

[14] developed an advanced deep-learning based detection algorithm named YOLOv8-MNC for identifying smoking behaviors in video frames, specifically, this model addresses very small and often occluded cigarette objects in realistic environments. This research uses the custom dataset model and achieves a mean average precision of (mAP@0.5) of 85.887% which represents a 5.7% improvement over the original YOLOv8 baseline. The authors also explain key challenges like light variations pose etc.

Face recognition has become a major research area due to the rapid growth of intelligent software applications including person recognition [15][16], healthcare [17][18], and e-learning [19], among others. However, reliable face identification remains a challenge because human facial features vary significantly under different conditions. [20] presented a comprehensive review of facial recognition and surveillance systems, covering underlying principles, key application domains, benchmark datasets, and feature extraction methods including appearance-based, model-based, and hybrid approaches. It also examines critical components such as feature selection, classification techniques, and evaluation protocols, while highlighting current challenges and future research directions for building more accurate and robust face recognition systems. [21] proposed a web-based application for detecting face in real time images and videos streams, comparing them with four pretrained algorithms i.e. Haar-Cascade, HOG-based frontal detector, MTCNN, and DNN and concluded that DNN and MTCNN performed better under different conditions. [22] proposed a fusion-based face recognition framework to improve the performance under challenging real-world conditions such as occlusions, pose variations, and lighting changes. The system integrated MTCNN for accurate face detection and alignment, FaceNet model for extracting discriminative facial embeddings, and an Adaptive Feature Fusion (AFF) module that combines identity features with contextual information such as pose, illumination, and mask presence. Their results demonstrated that the proposed attention and context-aware fusion mechanisms significantly improved the face recognition performance, therefore, this study was motivated to use MTCNN and FaceNet for face recognition in our research.

[23] proposed Hawk-Eye that is an AI-powered threat detection for Intelligent Surveillance Cameras, to integrate AI with real-time video surveillance to detect potential threats and abnormal behavior automatically. Surveillance systems using deep learning and computer vision have been widely studied for public safety applications such as violence detection, crowd monitoring, and anomaly recognition in video streams [24][25]. There is great demand for intelligent systems capable of automatically detecting anomalous events in videos, and a wide variety of approaches have been proposed to ensure public security [26]. However, despite the growing body of research, very limited work has been focused on policy-specific violations within academic institutions, particularly on the detection of cigarette smoking in university campuses, which remains a largely unexplored area in the surveillance literature.

The reviewed literature revealed that cigarette smoking is common among university students, even in medical and health-related institutes. It highlights the weak policy enforcement by institutes, significant research gap, and the need for real-time monitoring. Most of the studies have shown that DL models like YOLO, CNN, and YOLOv8 can accurately detect smoking behavior, even in challenging environmental conditions. However, the reviewed research strongly supports developing an enhanced AI-based cigarette smoking detection system for universities to improve monitoring and policy enforcement.

The aim of this research is to develop a surveillance system for cigarette smoking detection and student identification to promote a smoke-free environment on university campuses. The system captures visual evidence such as screenshots from CCTV cameras, identifies the student and sends quick notification to the authorized personnel. The system will keep **a** record of the number of smoking attempts by an individual and it will generate a penalty challan for individuals.

Novelty of the Study:

The novelty of this study is the development of integrated AI-powered surveillance application that combines cigarette smoking detection with automated student face identification and alert generation in a single system. Unlike other existing smoking detection systems that only focus on detection of smoking activities, the proposed approach provides a web dashboard that not only detects the cigarette smoking activity but also identifies the student and generates an email alert. Furthermore, the study introduces a custom cigarette smoking dataset collected from six publicly available datasets that improves detection accuracy.

Material and Methods:

This research aims to develop an AI system for cigarette detection, face recognition to recognize the individual's face and the web Dashboard for interaction. The proposed methodology for this system is shown in Figure 1. This section describes each phase of the system in detail.

The flow chart in Figure 2 shows the working of each component of the system. The cigarette detection process starts with the live video streams captured from the web camera, which is analyzed in real-time using DL, typically YOLOv8. This model is pretrained on custom datasets which contain images of both cigarette and non-cigarettes allowing it to accurately detect the presence of cigarettes, even when they appear in very small or partially obscured objects in a frame.

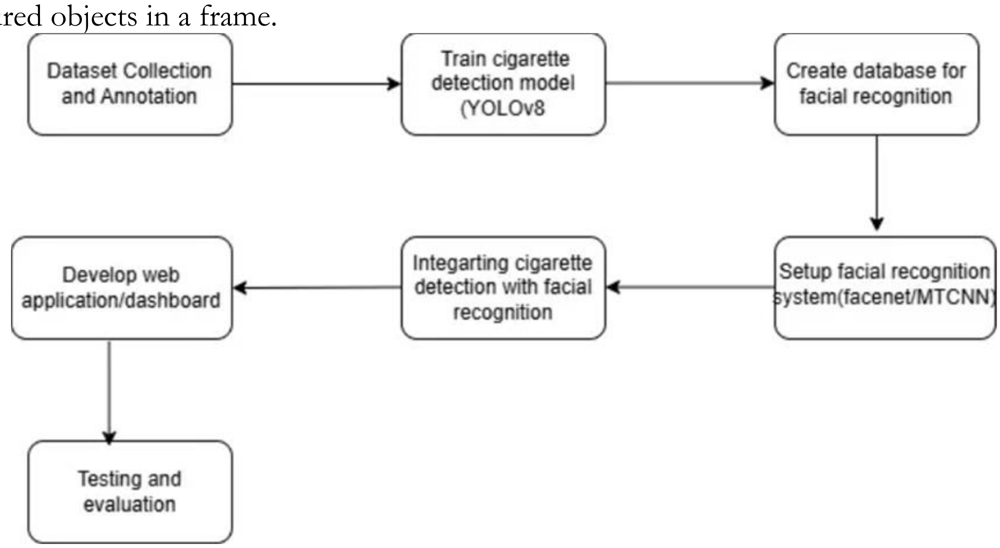

Figure 1. Methodology of the system

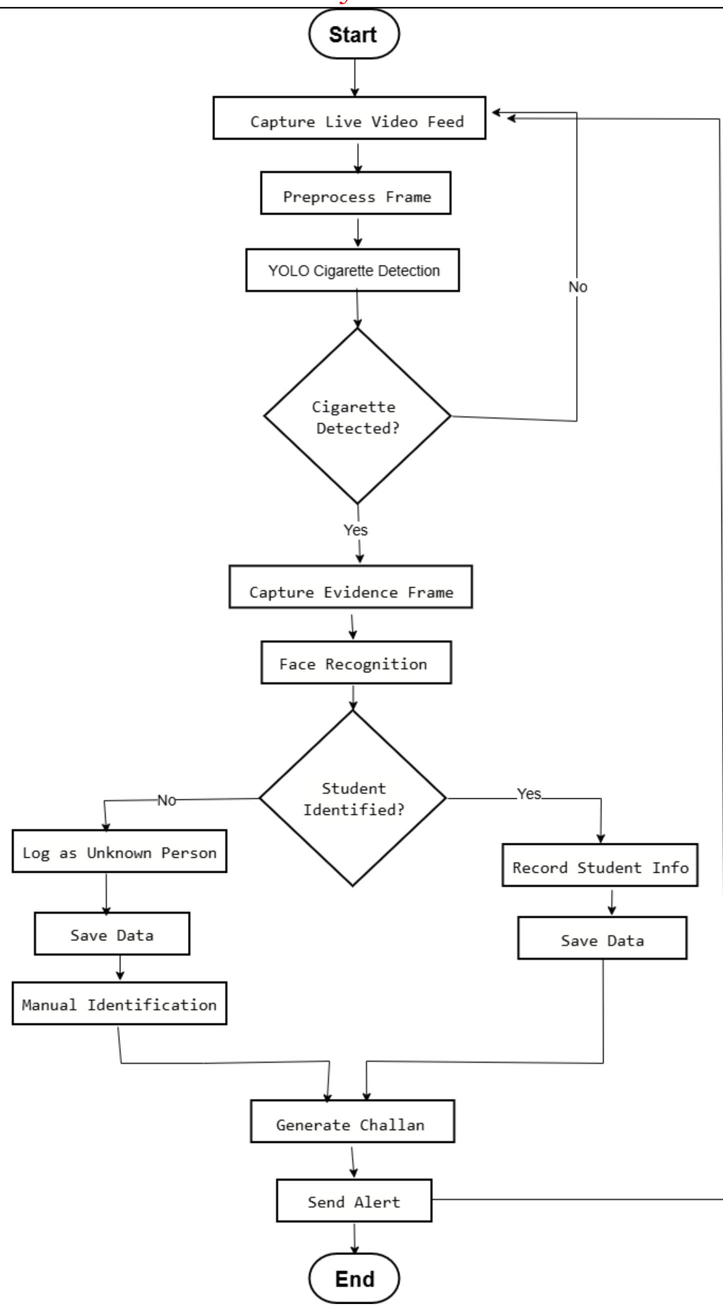

Figure 2. Flowchart of the system

When a cigarette is detected in a video stream, then the system saves the frame as evidence and prepares it for the next step in the workflow. Once the video frame with a detected cigarette is noticed, the system uses facial recognition techniques to identify the individual involved in cigarette smoking. This process starts with face detection, we implemented MTCNN, which locates and extracts facial regions from the frame. These facial frames are processed further using FaceNet for generating **a** unique face vector. It is a mathematical representation of the face. This vector is then compared against an already existing MongoDB database of the students' face vectors. If the face matches, the system then identifies the student and records their details along with the captured smoking activity. Otherwise, the person is tagged as unknown. The output on the dashboard includes evidence frames, identified names, timestamp, and related information. This dashboard allows authorized personnel to view activities, export reports, and monitor compliance in real-time, ensuring efficient detection and effective implementation of anti-smoke policies.

Proposed System Pipeline:

The proposed system operates as an end-to-end automated pipeline comprising five interconnected modules: data preparation, cigarette detection, face recognition, database matching, and alert notification. Those five modules are described in the following seven steps.

Step 1: Dataset Acquisition and Model Training:

The system pipeline starts with the acquisition and annotation of cigarette-smoking images collected from seven datasets [27][28][29][30][31][32][33] that are publicly available. These images are tagged using bounding box annotations and fed into the YOLOv8 model for training. YOLOv8 is chosen for its real-time object detection ability and high precision on custom sets of data. The model has been trained and validated repeatedly until sufficient performance parameters are achieved.

Step 2: Live Video Feed Acquisition and Preprocessing:

Once the system is deployed, it continuously receives video streams from the CCTV cameras. Each video frame is extracted and preprocessed by resizing and normalizing it to match the input specifications of the YOLOv8 model.

Step 3: Cigarette Detection (YOLOv8):

Each preprocessed frame is passed through the YOLOv8 model, which examines the frame to detect the presence of a cigarette. If no cigarette is present, the frame is disposed, and the system proceeds to scan the next frame in a continuous loop. If a cigarette is present and detected with a confidence score that is above the defined threshold (i.e. >0.6), the system then annotates the frame as a violation and proceeds to the next module. At this stage, a screenshot of the frame is captured with a timestamp to be stored as evidence for further processing.

Step 4: Face Detection and Recognition (MTCNN + FaceNet):

The captured frame is then passed to the facial recognition module, where MTCNN is used to scan and localize the face of the person smoking within the frame. MTCNN runs three cascaded stages (proposal, refinement, and output) to precisely detect face bounding boxes and facial features and landmarks. The detected and cropped face region is then sent over to FaceNet, which converts it into a 128-dimensional embedding vector that represents the person's unique facial features.

Step 5: Database Matching:

The produced FaceNet embedding is compared against a student face database that contains the facial data for 30 enrolled students. This matching is performed using cosine similarity. A similarity score above the threshold (>0.6) is considered as successful identification. The system then retrieves the student's basic information from database to send alerts. In case of no match being found, the person is flagged as an unknown individual. This may be due to either the person being unregistered or their face not being sufficiently clear. The recorded evidence is still saved, and the violation is marked for manual verification by relevant campus authorities.

Step 6: Challan Generation and Alert Notification:

After a successful identification, the system automatically generates and issues a disciplinary challan containing the student's registration details, the timestamp of the violation, along with the evidence screenshot being attached. This record is saved to the backend database and uploaded to the web dashboard at the same time, allowing for campus management to review all incidents and violations in real time. In addition, an automated email alert is also sent to the relevant disciplinary authority, enabling swift disciplinary action without the need for manual monitoring.

Step 7: Continuous Loop:

The system continues to process the live CCTV feed after an alert is generated and sent. This chain continues indefinitely and maintains uninterrupted surveillance across all the CCTV feeds. The entire cycle from frame capture to alert dispatch functions in real-time, which ensures timely detection and immediate action to smoking violations on the campus.

Datasets:

During the process of data collection, we initially used dataset from Mendeley Data and annotated it by ourselves, but later we went for dataset Cigarette Smoking [28][29][30][31][32][33] which we used to train our final model. There are some other datasets which we initially annotated and explored were:

Smokers' vs Non-Smoker:

This Dataset [27] contains total no of 24,00 raw images where (1200) images are of smoking category, and other (1200) images are of non-smoking category. This dataset contains various images such as cigarette smoking, smokers, people, coughing, taking inhalers, person on the phone, drinking water and much more.

The dataset was divided into approx. 84% training set (2016 images), 10% validation set (240 images) and 6% for testing (144 images).

Cigarette Detection:

This dataset is created from six datasets [28][29][30][31][32][33] that are publicly available on Rob flow. This Dataset has total 4,640 annotated images, in which 3,682 images are of smoking and remaining 958 are of non-smoking images accordingly. The 958 non-cigarette images contain similar objects as cigarette like pens, pencil and straws that may be considered as cigarette by the model. There was a problem of class imbalance with more cigarette images as compared to plain normal background images which help the model to focus more on the small object detection. The dataset was divided as approx. 75.2% into training set (3492 images), 17.2% validation set (796 images) and 7.6% for testing (352 images).

Student Database:

Generally, at the time of admission, information of each student along with picture is stored in the institutional database. That database can be used to match the individual who violated the smoking policy at campus. For the prototype, we have created a database of 30 enrolled students. For each student, basic information including name, roll no, contact details and facial images in controlled conditions were recorded. The collected facial images were processed through MTCNN for face detection and FaceNet for generating 128-dimensional embedding vector for each student which were stored in the database and serve as reference template for live face recognition during system operation.

Dataset Preprocessing:

The images were resized to 640x640 for further processing. This ensures that the images fed into the model are of the same size and helps in computational efficiency and learning of relevant features. After that, the standardization of orientation was performed to ensure all images are correctly oriented such that subjects and objects of interest like cigarettes are always presented in the same fashion as the model would see during real-world detection. Next is labeling with bounding boxes. Each image of this dataset has been annotated by drawing bounding boxes. More precisely, for every image in that dataset, a tightly bounding rectangle has been drawn around every cigarette appearing in the scene, including information about its location and size. These bounding boxes are highly relevant for the labeled data needed by YOLOv8 to learn how to recognize and locate cigarettes in new, unseen images. The annotation of the bounding box should be done with great precision, since it directly influences the model's detection performance. The careful pre-processing steps ensure that quality and consistency of data support the effective detection of cigarettes and thereby increase the accuracy and reliability of the model in practical deployment.

Model Training:

To detect cigarettes, the proposed system was based on the YOLOv8s (You Only Look Once, version 8 small) object detection model since it offers an ideal tradeoff between speed and accuracy particularly in real-time applications and localization of small objects such as cigarettes in diverse campus situations. The reason to choose YOLOv8 over other models like Faster R-CNN, RT-DETR, and EfficientDet was that it strikes the right balance between speed and accuracy for real-time surveillance. Faster R-CNN was somewhat slow for live CCTV processing, while newer models like YOLOv10 and YOLOv12 lacked mature tooling for custom dataset training at the time of development.

The training of YOLOv8 was performed on a dataset that has been carefully constructed, 75.2% of the pictures in total were used as training data (3492). In this process, the model was trained to localize and classify unseen images of cigarettes which form the essence of the detection of the surveillance solution.

Experimental Setup and Hyperparameters:

The model training was done on Google Colab's free tier, which provided access to NVIDIA Tesla K80 GPU with 12GB VRAM and 12 GB RAM, subject to session availability. The training was conducted using Python 3.10 with the Ultralytics YOLOv8 framework. The model was trained for 50 epochs with batch size 16, an initial learning rate of 0.01 and an early stopping patience of 10 epochs to prevent overfitting. Diverse values of augmentation (flipping, scaling, mosaic, and mixup) were used to ensure that the model would be adapted to the real-life conditions. The total training time was approximately 3 hours on free tier cloud GPU. Following the training the model achieved an inference speed of approximately 38ms per frame, equivalent to roughly 26 frames per second (FPS), which is sufficient for real-time CCTV based surveillance. After the model training, the YOLOv8s was then subjected to the validation process, and the final model performance is measured on the unseen test set.

Predictions on Live Video:

The trained YOLOv8s model can process live camera feeds in real-time, and the incoming image is checked against the potential cigarettes, which demonstrates that the performance was rather strong and can be used for the implementation of the system. The model gives confidence score of each cigarette denoted in the frame that is the probability of the object that is depicted in the frame to be a cigarette. Such real-time feedback will help the administrators to get the performance of the detection and respond to it in a timely and evidence-based manner. It allows the system to monitor the development of campus in a precise and transparent manner since it provides reports on the position, as well as the confidence of each prediction.

Face Recognition Module:

When a cigarette is detected in one of the video frames, it is said to trigger the face recognition feature of the system. Multi-task cascaded convolutional neural network (MTCNN) is a popular deep learning framework for face recognition.

MTCNN:

MTCNN operates in three stages. Initially, the photo is scanned to propose possible face areas; secondly, the proposals are narrowed down to eliminate false positives, as well as to better define the facial area; third, it searches the face to detect significant facial features, such as eyes, nose and mouth. This ensures that the cropped and oriented face is identified that will go for the next step.

FaceNet:

FaceNet is an opensource facial recognition system developed by Google. Once MTCNN isolates and aligns the face, our system uses FaceNet for feature extraction and recognition. It processes the detected face and converts it into a unique numerical representation known as face embedding. This embedding captures the distinctive features of a person's face in a way that remains consistent even under different lighting conditions, angles, or expressions. When a smoking violation is detected, the system compares this embedding against the stored face vectors of enrolled students in the MongoDB database. If a close match is found, the student is identified and their information is retrieved. If no match is found, the individual is logged as unknown and flagged for manual review.

Web Dashboard:

NoSmoke web dashboard was developed for real-time monitoring and enforcing the no-smoking policies on the campus. Manual surveillance is not always fair and takes a lot of time and efforts. It is prone to error, hence does not suit large or open areas. The overall objective of the dashboard design was to provide campus authorities with centralized access which would be convenient in tracking the cases of smoking, the case perpetrators, documentation and appropriate time management of their activities. To carry out database operations, the dashboard was developed using Fast API as the frontend and backend and connected to MongoDB database. The system leverages modern technology stack with React.js for building a dynamic, component-based dashboard enhanced by Tailwind CSS for efficient styling and WebSocket for real-time two-way communication to stream live video and detection updates. On backend, python provides a versatile foundation for data processing, machine learning, image manipulation using libraries such as OpenCV, while FastAPI serves as a high-performance web framework to seamlessly integrate the frontend with the backend's detection and recognition services. All communications are secured, and only the administrators can access the account. Everything including record keeping, enforcement and generation of reports is automated that enables the administration of no-smoking policies to be done in real-time on campuses. The dashboard is provided with simple functions including:

Home Page Navigation:

It has two main buttons, Live Detection and logs, one can instantly choose to enable live monitoring or to review past detections. It also provides a menu bar through which one can move to other pages such as History, Database, Challan and Admin.

Live Detection:

It allows the user to do real time cigarette detection and built-in face recognition with the use of the web camera. Once one frame is captured, the system scans that frame. If cigarette is detected it tries to identify the person. On detecting students, a challan is drawn (fine notification) and an email dispatched to the concerned student. Figure 3 shows the Live Detection interface of the NoSmoke dashboard, where real-time cigarette detection alerts and video feed are displayed.

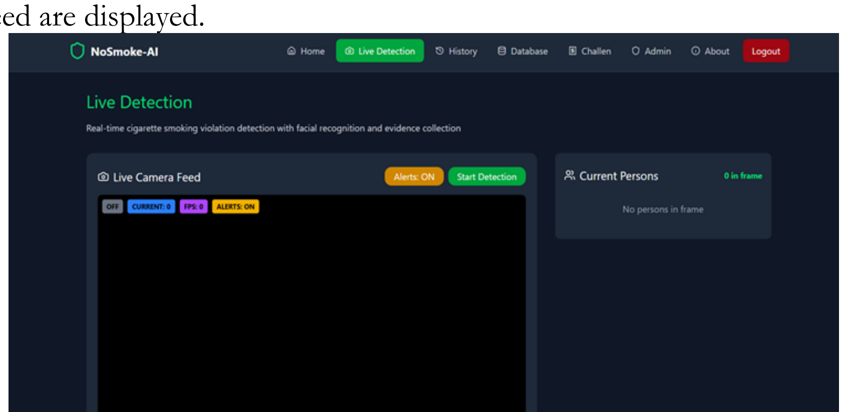

Figure 3. Live Detection for NoSmoke Dashboard

Historical logs (History):

Displays a historical list of all the previous instances of detection with screen shots of the live feed of the camera. All the information regarding each smoking violation is presented with the identity of the detected person, the moment of the accident, and the confidence rating along with an alert button for notifying the administration. It assists administrators in sorting the logs, exporting evidence and survey system activity.

Database:

This interface allows student information to be entered into and updated in the database to be used for face recognition. Once the cigarette is detected, the face recognition system matches the person's face with the student database. Figure 4 shows the Student Database section of the NoSmoke Dashboard, where student information is manually inserted, updated, or deleted, and used for identification during detection events.

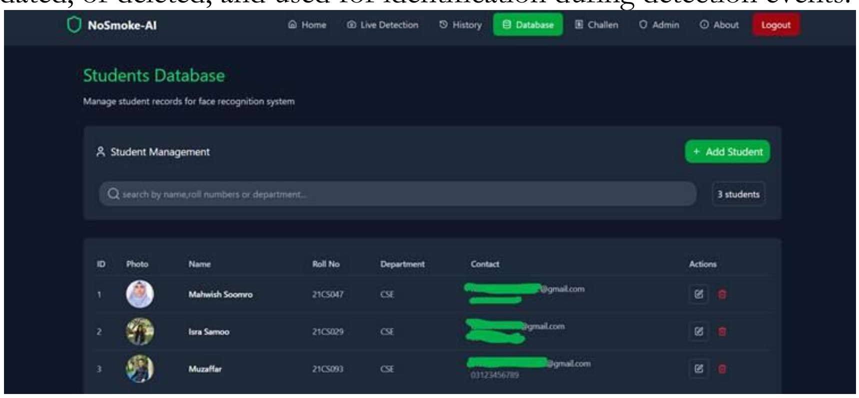

Figure4. Student Database for NoSmoke Dashboard

Challan Information:

This page shows the list of all challans generated for smoking violations. It tracks the status of all the challans that indicate whether they have been paid. It also provides actionable options to the admins to update the payment status and manage the enforcement process. Figure 5 shows the Challan History page of the Dashboard, displaying all issued fines along with their payment status and options to mark them as paid or download the details.

Admin Access and Security:

Limits access of the system to authorized users using administrator credentials which enhances the level of data security and governance. Administrators are allowed to add new students, look at detection logs, manage the challans and control the system settings to ensure that the sensitive information and administrative controls are secure.

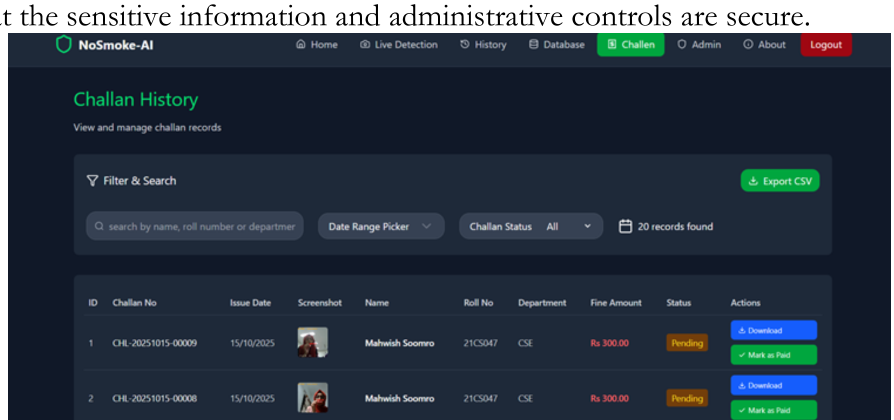

Figure 5. Challan History for NoSmoke Dashboard

Result and Discussion:

Cigarette Detection:

The evaluation of **the cigarette smoking detection model**is performed using two datasets. In this section, we have presented key metrics like confusion matrices, precision, recall, F1 score and accuracy to evaluate the model's performance on the two datasets in comparison with each other.

Dataset 1 (Smoker Vs Non-Smoker):

The Dataset contains 2400 images consisting of both smoking and non-smoking categories. This dataset contains many different actions performed by people such as coughing and drinking, allowing the model to learn to distinguish cigarettes from similar objects. The confusion matrix in Figure 6 shows the model predictions for each class. The diagonal elements represent the number of correctly predicted cases for each class, whereas off-diagonal elements represent misclassified instances. Thus, the higher values on the diagonal represents better model performance as it indicates the number of correctly classified samples. The model performs well in detecting cigarettes, but there are still a large number of false positives which need to be reduced. The model has accuracy of 85.0%.

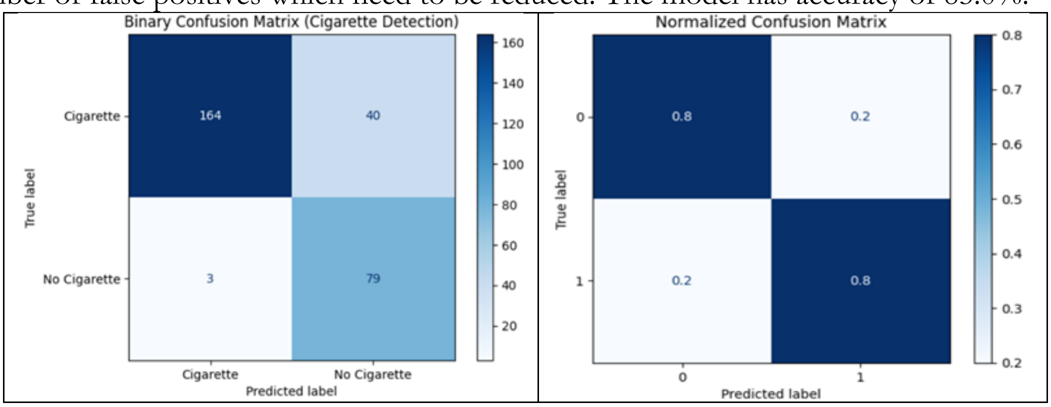

Figure 6. Confusion and Normalized Confusion matrix of Dataset 1

Figure 7 shows training graphs of the performance of model during training. Visualizations support the tracking of metrics for box loss, classification loss and precision recall in graph form by epoch, showing how these things improved over time.

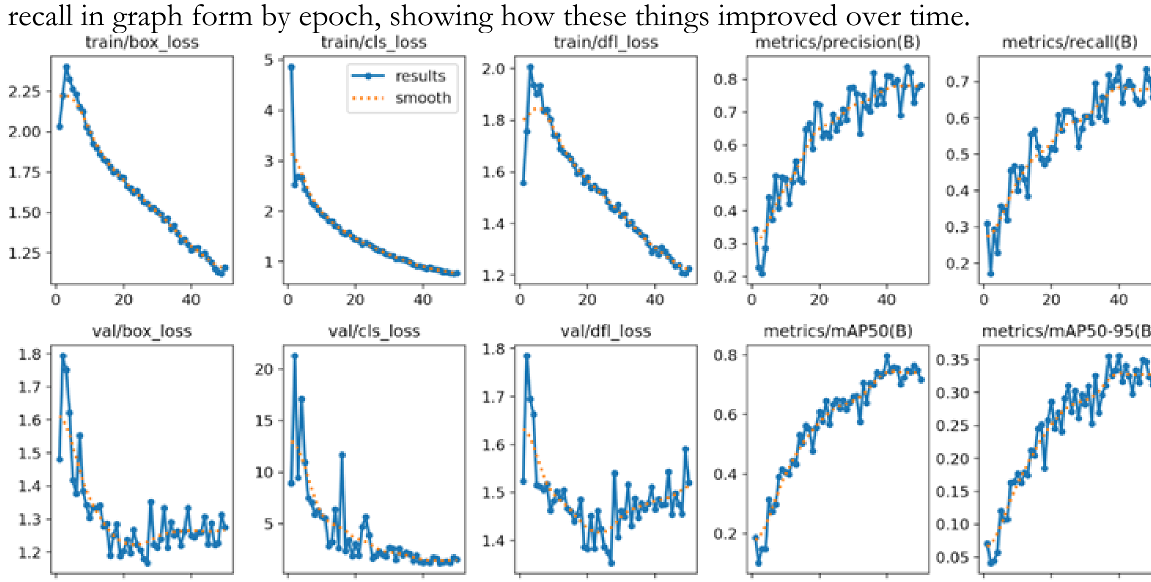

Figure 7. Training and validation curves of Dataset 1 (Smoker vs Non-Smoker)

Dataset 2 (Cigarette Smoking):

This dataset contains images of both smoking and non-smoking categories, including many similar objects that look like cigarettes, such as straws, pens, and pencils.This enables the model to more accurately distinguish cigarettes from similar objects. This dataset was used for model training on two different YOLO versions, i.e. YOLOv8s and YOLO11s. The YOLOv8s model, with an accuracy of 90.0%, was selected as the final detection model to be integrated into the No Smoke AI system due to its higher accuracy than the other trained models. YOLOv11 model had 88.6% accuracy for this dataset. Despite the faster inference of newer versions of YOLO, YOLOv8 is still preferred for its accuracy on smaller objects. Confusion matrix in Figure 8 shows the model predictions for each class.

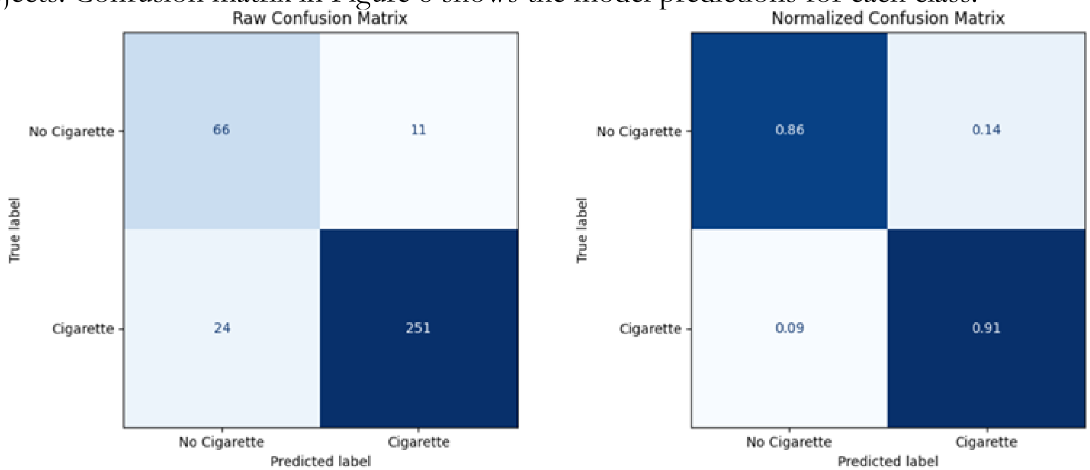

Figure 8. Confusion matrix of Dataset 2 using YOLOv8

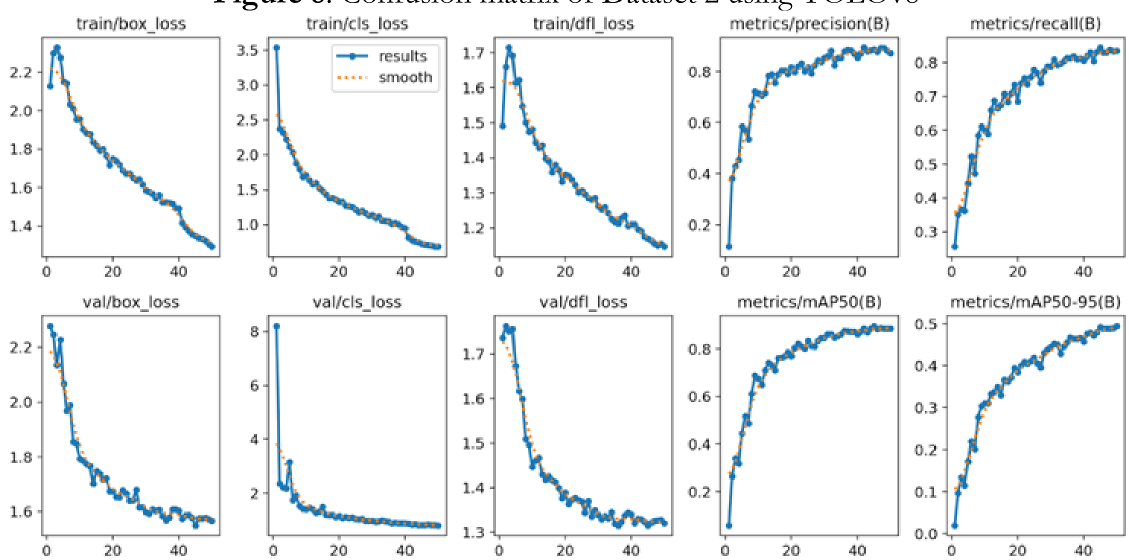

Figure 9. Training and validation curves of Dataset 2 using YOLOv8

The plots in Figure 9 indicate that YOLOv8 training is performing well in general, with some room to work on stability in validation. The first row shows that the train box loss, classification loss, and Distribution Focal Loss (DFL) loss decrease steadily with the number of epochs, indicating that the model is becoming more accurate in enclosing tight boxes around cigarettes, labeling them correctly, and estimating accurate box boundaries. Meanwhile, training precision and recall curves continue to increase, and thus the model not only is generating fewer false positives (precision up) but also is detecting more true objects (recall up) as it continues to train. The validation box and classification losses in the second row are sharp at the beginning and then becomes flat and slightly spiky, indicating that the model is quickly learning the major patterns but there is no serious upward trend in validation performance, likely because there were relatively few or highly diverse validation samples; yet, they continue to decrease, so there is no evidence of overfitting in these 50 epochs. The validation DFL loss also reduces initially and then swings around a relatively stable value which indicates that the quality of box boundaries on unknown images is not worsening. These two graphs, validation mAP@0.5 and mAP@0.5:0.95 are the most important, as they both increase consistently with epochs, i.e. the detector is achieving progressively higher accuracy on average across the Intersection over Union (IoU) thresholds, and they are also increasing steadily, i.e. training can possibly be continued a little longer or you can save a late-epoch checkpoint in case your model has reached the best possible performance. Generally, the overall trends of these curves indicate that your training configuration (learning rate, augmentations, batch size) works, the model is learning meaningful features, and the most obvious next steps would be to reduce validation noise (e.g. more data or better balance) and maybe run a few extra epochs with early stopping on the best mAP.

The dataset 2 was also trained using YOLO11. Confusion matrix in Figure 10 shows the model predictions for each class.

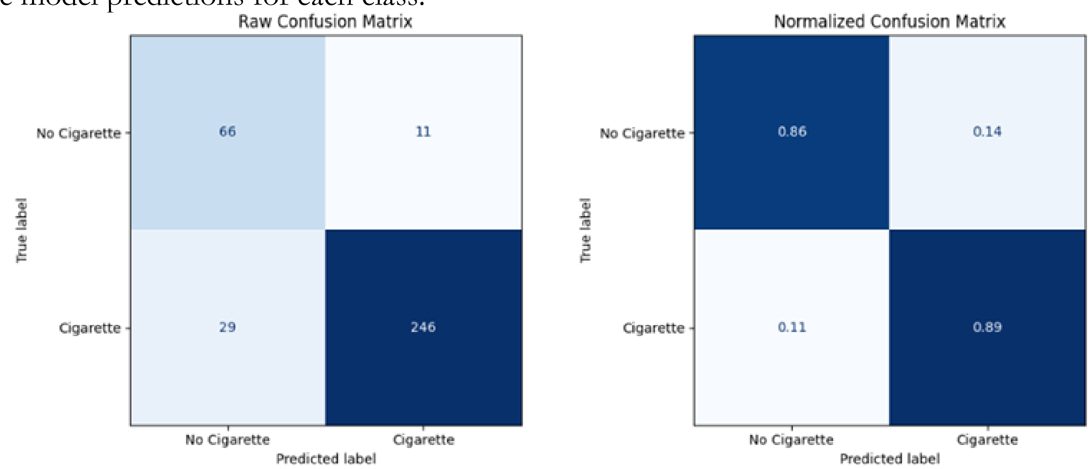

Figure 10. Confusion and Normalized Confusion matrix of Dataset 2 for YOLO11

Figure 11 shows summary of the performance of model during training. Visualizations support the tracking of metrics for box loss, classification loss and precision recall in graph form by epoch, showing how these things improved over time. Both the graphs of Figure 9 and Figure 11 indicate that the proposed model trains in a stable and effective manner, but its eventual performance remains slightly lower than the result of YOLOv8 on the same dataset. Trends of the training box, classification, and DFL losses reduce gradually, and accuracy and recall grow in a steady fashion, which proves that the model is learning to draw right boxes, select the correct class, and can see the majority of target objects. On the validation side, the losses decrease rapidly and after that show some minor fluctuations and this is a good sign that there is good generalization and no evident overfitting. But a direct comparison with the best YOLOv8s run shows that the eventual validation mAP at 0.5, and mAP at 0.5:0.95 here are slightly lower, and the validation losses are not such low values that means that YOLOv8 can still be expected to reach slightly higher overall detection accuracy on this dataset. Simply stated, this version is good and well-trained, although YOLOv8s is the stronger one in this particular task.

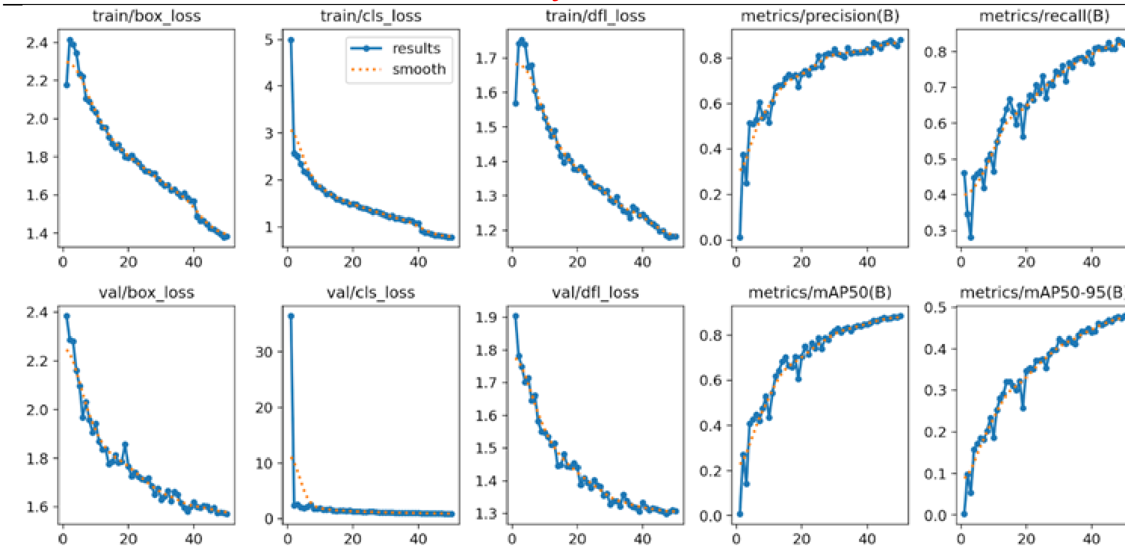

Figure 11. Training and validation curves of Dataset 2 (Cigarette Smoking) using YOLO11

To develop a final model, we chose the Cigarette Smoking dataset collected by various open sources and narrowed down with the help of Roboflow. The choice was made because it is a very specific dataset that directly dealt with cigarette-related images and had both smoking and non-smoking scenes that were well-annotated. Notably, the non-smoking pictures purposely used objects like pens, pencils, and straws that can easily be mistaken with cigarettes to condition the model to distinguish cigarettes with other objects that appear in a similar manner. Its annotated images were selected in a more realistic set of surveillance situations with a mixture of angles, distances and the type of visual clutter experienced in real CCTV surveillance. This had the benefit of making the trained model reliable in detecting cigarette smoking in the real world. Moreover, the increased amount of cigarette images in this data set made sure that the model focuses on the detection task. The Cigarette Smoking dataset therefore offered a far more promising platform to precise, real-time cigarette identification and more closely match the goals and requirements of our campus surveillance initiative. Figure 12 shows the screenshot captured when the cigarette was detected in the live camera feed of the NoSmoke dashboard.

Figure 12. Screenshot captured after Cigarette Detection

Face Recognition:

The results of the face recognition module are presented in this section. The face recognition module employs pretrained MTCNN and FaceNet models rather than a trained-from-scratch classifier, and its role in the system is to match the detected face against a pre-built student database of 30 enrolled students using cosine similarity. For the experimental results, 6 images per 30-student database have been taken, out of which 14 were misclassified. The experimental results have shown satisfactory performance over all evaluation metrics, achieving an overall accuracy of 92.2%, precision of 91.8%, recall of 92.2%, and F1-score of 91.9%. The false acceptance rate was recorded at 3.1%, means the system rarely misidentified one student as another. The false rejection rate stood at 7.8%, that is primarily because of poor lighting or partial face occlusion reduced the quality of the captured face.

Figure 13 shows the results for the person identified if the information of student is available in the database. Figure. 14 shows the result for an unknown person considering that the person's information is not available in the database.

Figure 13. Student Identified in NoSmoke system

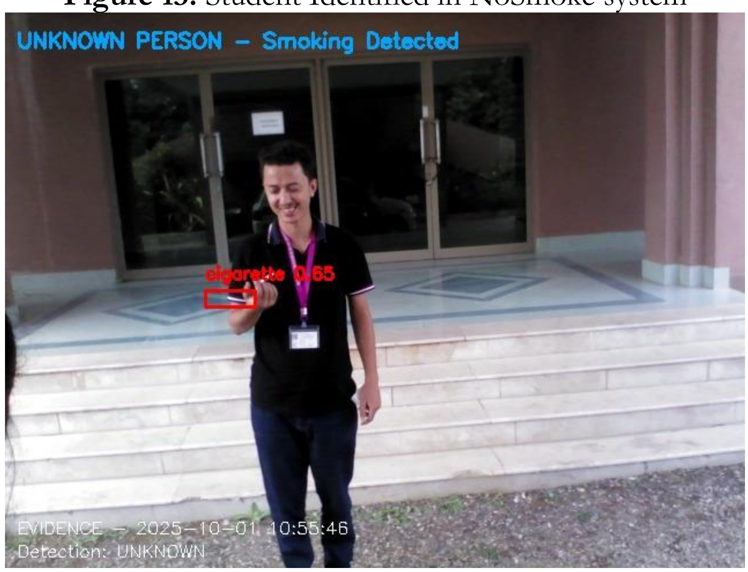

Figure 14. Unknown person not identified in NoSmoke system

Email Alerts and Challan Generation:

The NoSmoke system is intended to react in real time and automatically when any case of violation of smoking regulations is observed using the live camera image. Upon occurrence of such an incident, such as when a student is found smoking cigarettes, the system immediately creates an email notification that will include the message to the individual notifying them of the violation that was detected and providing details associated with it, including time, date, location, etc. The automated email notification sent to a student after a system violation has been recorded is shown in Figure 15.

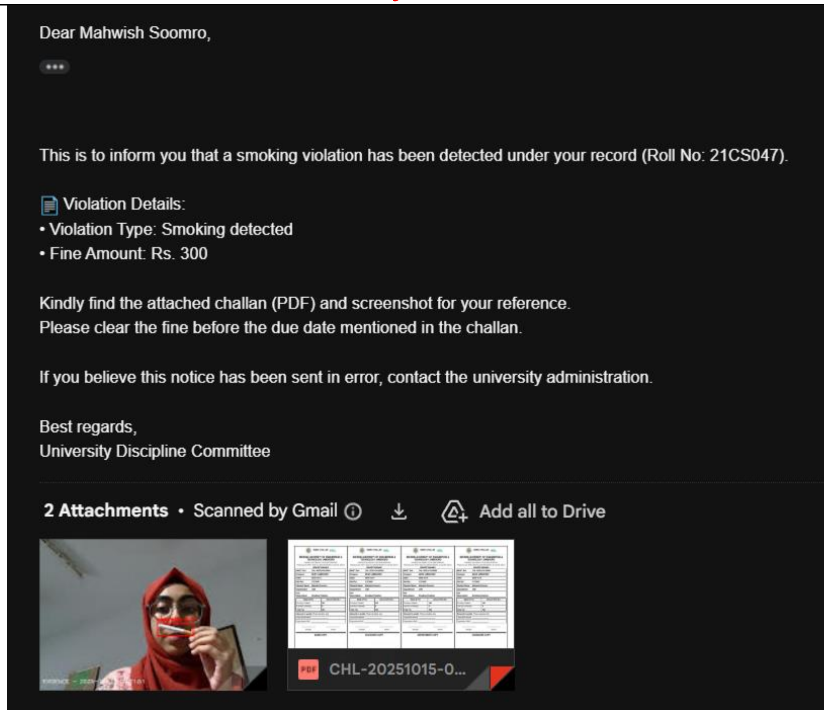

Figure 15. Email alert sent upon violation

Besides email alerts, the system generates a challan (fine) per confirmed smoking case. A student on recognition is added to the database, which connects the violation and the record of the student, and the information is marked as to be reviewed by an administration. The information about the challan (paid/unpaid) can also be found on a challan page on the web dashboard where its status can be updated and controlled by the authorities.

Comparative Analysis:

The proposed system is compared with existing literature on cigarette detection as shown in Table 1. The comparison is based on method used, dataset, detection accuracy, and features. Although we acknowledge that the performance of the proposed system is somewhat lower, but it provides an end-to-end pipeline of cigarette detection, face recognition and alert system, significantly designed for the academic institution setting.

Table 1. comparison of the proposed system with existing studies on cigarette detection.

Study	Method	Dataset	Accuracy/mAP	Face Recognition	Alert System	Student Identification
[7]	CNN, SVM, KNN, Decision Tree, Gradient Boosting	Binary smoking classification dataset	CNN performed 100%	No	No	No
[9]	Single Shot Detector (SSD)	Custom no smoking images in public area	97.5%	No	Yes	No
[10]	Lightweight custom CNN	ICSSD indoor worksite	mAP50 0.917	No	Yes	No

		dataset				
[12]	YOLOv3	Custom cigarette images	92.5%	No	No	No
Proposed System	**Yolov8, MTCNN, FaceNet**	**Custom cigarette dataset**	**90%**	**Yes**	**Yes**	**Yes**

System Evaluation:

A controlled test was conducted to evaluate the email alert mechanism, in which 50 simulated smoking violation events were recorded across different network conditions. For each event the system was expected to generate an email alert to the university authorities and student.

Table 2. Notification delivery evaluation

Test Scenario	Events Triggered	Email Delivered	Delivery Rate
Normal network condition	20	20	100%
Moderate network load	15	14	93.3%
High network load	10	9	90.0%
Multiple simultaneous alerts	5	4	80%
Total	**50**	**47**	**94.0%**

The web dashboard, automated challan generation, email notification modules were evaluated for the practical reliability of the proposed system. Table 2 shows that an overall 94% success rate was achieved for email delivery across all test conditions. The 6% failure rate was due to multiple simultaneous alerts and network congestion. Table 3 shows the web dashboard functionality evaluation on four functional criteria through structured testing. Each criterion was tested across 20 sessions and rated on successful execution. The results showed all the dashboard features performed reliably. The success rate of challan generation was 90% due to rendering delays when multiple challans were generated.

Table 3. Dashboard functionality evaluation

Dashboard Feature	Tests Conducted	Successful Execution	Success Rate
Real-time violation log display	20	19	95%
Evidence screenshot rendering	20	20	100%
Student record retrieval	20	20	100%
Challan PDF generation	20	18	90%

To validate the robustness of the reported results, the experiments were repeated five times on the same test set and key performance metrics were recorded for each run. The system was able to detect cigarettes most of the time and the detection accuracy was around 90.1%± 0.23% across all runs. The face recognition accuracy was around 92.2%±0.3% for each run, with a 95% confidence interval of [91.83%, 92.57%]. The narrow standard deviation was observed across all metrics. This validates that the results of the proposed system are stable and reproducible and can be used for real-time surveillance of academic environment.

Conclusion:

This research proposed an end-to-end AI-based system that monitors cigarette smoking on campus and identifies students using computer vision and DL. The YOLOv8 is used to precisely identify cigarettes from live video frames, in even challenging conditions of surveillance cameras. To enhance the detection accuracy of the model, the dataset is first preprocessed using image resizing, correcting image orientation, and robust annotation. Once the cigarette is detected, the system then uses MTCNN and FaceNet to identify the face of student. Moreover, the system includes a web dashboard to monitor live, review historical records, and generate automated email and challan (fine). This research proposes that AI and ML can significantly boost campus safety programs and offer a platform of scalable intelligent surveillance in the community or institutional environment. This saves manual labor of campus staff and offers an efficient problem-free solution to maintaining a smoke-free campus.

Despite the promising results of the proposed system, there are still some limitations that must be acknowledged. Although the face recognition module performs well, in some scenarios where face is partially occluded with bag, hands or face mask, the detection accuracy may decrease. In crowded areas such as canteen, corridors, and campus entrances where multiple individuals may appear in a frame, the proposed system can detect only one individual in a frame and does not support tracking multiple persons smoking in one frame. The deployment of AI powered surveillance in campus often raises privacy and ethical concerns, that require informed consent and secure storage protocols for the storage of facial embeddings.

Acknowledgement:

The authors would like to express their sincere gratitude to the Department of Computer systems, Mehran University of Engineering and Technology, Jamshoro, for assisting in this research.

Author's Contribution:

Zartasha Baloch: Apart from writing the manuscript, she supervised the entire research, providing expert guidance on research methodology, data analysis, and ensuring scientific accuracy.

Irfan Ali Bhacho: Contributed to manuscript writing and NoSmoke web app development.

MadehaMemon: Contributed to manuscript writing and data analysis.

Mahwish Soomro: Contributed to designing, and developing the NoSmoke web app, dataset acquisition & preprocessing, training and testing DL models, and interpreting results.

Isra Samo: Contributed to the design and graphics of the dashboard to make them visually appealing.

Muzaffar Hussain: Created student database, managed on-field research, capturing cigarette images and analyzing app performance in real-world settings.

Conflict of interest: The authors declare that there is no conflict of interest.

References:

[1] Michelle Ren, Shahrdad Lotfipour, "Nicotine Gateway Effects on Adolescent Substance Use," *West. J. Emerg. Med.*, vol. 20, no. 5, pp. 696–709, 2019, doi: 10.5811/westjem.2019.7.41661.

[2] T. N. Kozhevnikova *et al.*, "The problem of teenage tobacco smoking: yesterday, today, tomorrow," *Pediatr. Cons. Medicum*, vol. 0, no. 2, pp. 101–108, Jun. 2021, doi: 10.26442/26586630.2021.2.200994.

[3] Sang hee Park, "Smoking and adolescent health," *Korean J. Pediatr.*, vol. 54, no. 10, 2011, doi: 10.3345/kjp.2011.54.10.401.

[4] Mishary Alhindal, Jood Janahi, "Impact of smoking on cardiovascular health: Mechanisms, epidemiology and specific concerns regarding congenital heart disease," *Int. J. Cardiol. Congenit. Hear. Dis.*, vol. 20, p. 100581, 2025, doi: https://doi.org/10.1016/j.ijcchd.2025.100581.

[5] "Addiction to Tobacco Smoking and Vaping - PubMed." Accessed: Jul. 05, 2026. [Online]. Available: https://pubmed.ncbi.nlm.nih.gov/37441760/

[6] Emine BEYAZ, Sonay GÖKÇEOĞLU, "Level of Smoking Addiction and Affecting

Factors in Nursing Students," *Namık Kemal Tıp Derg.*, vol. 10, no. 3, pp. 325–331, 2022, doi: 10.4274/nkmj.galenos.2022.45220.
[7] Muhammad Fahrurrozi, Germanus Naru, "Optimizing Machine Learning Algorithms to Accelerate Smoking Detection," *Int. J. Informatics Eng. Comput.*, vol. 1, no. 2, pp. 38–55, 2024, doi: 10.70687/ijimatic.v1.i2.42.
[8] Mirko Casu, Francesco Guarnera, Giusy Rita Maria La Rosa, "Smoking Detection and Cessation: An Updated Scoping Review of Digital and Mobile Health Technologies," *IEEE J. Biomed. Heal. Informatics*, 2025, doi: 10.1109/JBHI.2025.3549255.
[9] A. Riansyah, A. T. Yulianto, N. A. Iahad, M. Qomaruddin, M. Taufik, and L. K. Wardhani, "Implementation of Single Shot Detector Method: Smoking Detection in No Smoking Area," pp. 267–274, Dec. 2025, doi: 10.1109/EECSI67060.2025.11290283.
[10] Meng Wang, Mei Li, "CASIA-Net: An indoor work site smoking detection framework," *Comput. Ind.*, vol. 173, p. 104383, 2025, doi: https://doi.org/10.1016/j.compind.2025.104383.
[11] "Face Authentication using MTCNN and FaceNet | Semantic Scholar." Accessed: Jul. 05, 2026. [Online]. Available: https://www.semanticscholar.org/paper/Face-Authentication-using-MTCNN-and-FaceNet-Pandit-Nikalje/a348606480e20dabf3a0173a4dde357fcb450851
[12] "Real Time Cigarette Detection using Deep Learning Models." Accessed: Jul. 05, 2026. [Online]. Available: https://www.researchgate.net/publication/371171533_Real_Time_Cigarette_Detection_using_Deep_Learning_Models
[13] Ali Khan, Mohammed A.M. Elhassan, "Deep learning-based smoker classification and detection: An overview and evaluation," *Expert Syst. Appl.*, vol. 267, 2025, [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0957417424030756
[14] Zhong Wang, Lanfang Lei, "Smoking behavior detection algorithm based on YOLOv8-MNC," *Front. Comput. Neurosci.*, vol. 17, 2023, doi: https://doi.org/10.3389/fncom.2023.1243779.
[15] Hang Du, Hailin Shi, Dan Zeng, "The Elements of End-to-end Deep Face Recognition: A Survey of Recent Advances," *arXiv:2009.13290*, 2021, [Online]. Available: https://arxiv.org/abs/2009.13290
[16] I. Masi, Y. Wu, T. Hassner, and P. Natarajan, "Deep Face Recognition: A Survey," *Proc. - 31st Conf. Graph. Patterns Images, SIBGRAPI 2018*, pp. 471–478, Jul. 2018, doi: 10.1109/SIBGRAPI.2018.00067.
[17] C. Bisogni, A. Castiglione, S. Hossain, F. Narducci, and S. Umer, "Impact of Deep Learning Approaches on Facial Expression Recognition in Healthcare Industries," *IEEE Trans. Ind. Informatics*, vol. 18, no. 8, pp. 5619–5627, Aug. 2022, doi: 10.1109/TII.2022.3141400.
[18] L. F. Gomez, A. Morales, J. R. Orozco-Arroyave, R. Daza, and J. Fierrez, "Improving parkinson detection using dynamic features from evoked expressions in video," *IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit. Work.*, pp. 1562–1570, Jun. 2021, doi: 10.1109/CVPRW53098.2021.00172.
[19] J. O.-G. Roberto Daza, Luis F. Gomez, Aythami Morales, Julian Fierrez, Ruben Tolosana, Ruth Cobos, "MATT: Multimodal Attention Level Estimation for e-learning Platforms," *arXiv:2301.09174*, 2023, [Online]. Available: https://arxiv.org/abs/2301.09174
[20] Aili Wang, Haibin Wu, "Advances in Face Recognition: A Comprehensive Review of Feature Extraction and Dataset Evaluation," *Electronics*, vol. 15, no. 2, p. 338, 2026,

doi: https://doi.org/10.3390/electronics15020338.

[21] "A web-based application for face detection in real-time images and videos - IOPscience." Accessed: Jul. 05, 2026. [Online]. Available: https://iopscience.iop.org/article/10.1088/1742-6596/2161/1/012071

[22] Loris Nanni, Hamed Taherdoost, "Combining MTCNN and Enhanced FaceNet with Adaptive Feature Fusion for Robust Face Recognition," *Technologies*, vol. 13, no. 10, p. 450, 2025, doi: https://doi.org/10.3390/technologies13100450.

[23] "Hawk-Eye: An AI-Powered Threat Detector for Intelligent Surveillance Cameras | IEEE Journals & Magazine | IEEE Xplore." Accessed: Jul. 05, 2026. [Online]. Available: https://ieeexplore.ieee.org/document/9408578

[24] E. Abdulwahab and A. K. Younis, "Surveillance System Based on Face Recognition Utilizing Deep Learning Technique," *1st Int. Conf. Emerg. Technol. Dependable Internet Things, ICETI 2024*, 2024, doi: 10.1109/ICETI63946.2024.10777082.

[25] A. Hamza, Z. H. But, U. Arif, Samiya, M. A. Asad, and M. Naeem, "Autonomous AI Surveillance: Multimodal Deep Learning for Cognitive and Behavioral Monitoring," Jul. 2025, Accessed: Jul. 05, 2026. [Online]. Available: http://arxiv.org/abs/2507.01590

[26] R. Thangaraj, A. Kishore Abinash, S. Kaviya, and R. Monika, "Deep Learning Enabled Traffic Surveillance and Police Verification Platform," *2026 IEEE 15th Int. Conf. Commun. Syst. Netw. Technol. CSNT 2026*, pp. 283–290, 2026, doi: 10.1109/CSNT69054.2026.11502043.

[27] A. Khan, "Dataset containing smoking and not-smoking images (smoker vs non-smoker)," vol. 1, 2020, doi: 10.17632/7B52HHZS3R.1.

[28] "Top Cigarette Computer Vision Models | Roboflow Universe." Accessed: Jul. 05, 2026. [Online]. Available: https://universe.roboflow.com/search?q=class%3Acigarette+trained+model

[29] "cigarette Object Detection Dataset by Eurus CV." Accessed: Jul. 05, 2026. [Online]. Available: https://universe.roboflow.com/eurus-cv/cigarette-zoio9

[30] "Cigarette Object Detection Dataset by beyhan." Accessed: Jul. 05, 2026. [Online]. Available: https://universe.roboflow.com/beyhan/cigarette-zf1g5

[31] K. M. F. James, K. Heiwolt, D. J. Sargent, and G. Cielniak, "Lincoln's Annotated Spatio-Temporal Strawberry Dataset (LAST-Straw)," *Lect. Notes Comput. Sci.*, vol. 15625 LNCS, pp. 46–63, 2025, doi: 10.1007/978-3-031-91835-3_4/SAVE-RESEARCH.

[32] P. Werukanjana, N. Permpool, and P. Sa-Nga-Ngam, "Integration of Object Detection in Crime Scene Investigation," *Eng. Access*, vol. 11, no. 1, pp. 13–23, Jan. 2025, doi: 10.14456/MIJET.2025.2.

[33] "pen - v1 pen-version2." Accessed: Jul. 05, 2026. [Online]. Available: https://universe.roboflow.com/yolov5-h3oo1/pen-zyj9j/dataset/1