

Multi Streams Face Recognition System for Student Tracking

BY

CHOW MUN KENT

A REPORT

SUBMITTED TO

Universiti Tunku Abdul Rahman

in partial fulfillment of the requirements

for the degree of

BACHELOR OF COMPUTER SCIENCE (HONOURS)

Faculty of Information and Communication Technology

(Kampar Campus)

OCTOBER 2025

COPYRIGHT STATEMENT

© 2025 Chow Mun Kent. All rights reserved.

This Final Year Project proposal is submitted in partial fulfillment of the requirements for the degree of **Bachelor of Computer Science (Honours)** at Universiti Tunku Abdul Rahman (UTAR). This Final Year Project proposal represents the work of the author, except where due acknowledgment has been made in the text. No part of this Final Year Project proposal may be reproduced, stored, or transmitted in any form or by any means, whether electronic, mechanical, photocopying, recording, or otherwise, without the prior written permission of the author or UTAR, in accordance with UTAR's Intellectual Property Policy.

ACKNOWLEDGEMENTS

I would like to express my sincere thanks and appreciation to my supervisors, Dr Aun Yichiet and moderator Dr Chai Tong Yuen who has given me this bright opportunity to engage in Multi Streams Face Recognition System for Student Tracking. It is my first step to establish a career in this research and development process. A million thanks to you.

To a very special person in my life, Kenny Roger, for his patience, unconditional support, and love, and for standing by my side during hard times. Finally, I must say thanks to my parents and my family for their love, support, and continuous encouragement throughout the course.

ABSTRACT

With the increasing demand for efficient space utilization and data-driven management in academic libraries, accurate and real-time occupancy monitoring has become an essential requirement. Traditional manual people counting methods are labour-intensive, error-prone, and incapable of providing continuous or historical insights for operational planning. To address these limitations, this project proposes an intelligent computer vision-based people counting and occupancy analytics system specifically designed for library environments.

The proposed system utilizes a real-time video processing pipeline that integrates YOLOv8 for human detection with ByteTrack for multi-object tracking. A strict zone-transition-based counting mechanism is implemented to accurately detect entry and exit events while minimizing double counting and false detections. Multiple anti-duplicate strategies, including spatial suppression, temporal cooldown, and minimum track age validation, are applied to improve robustness under real-world conditions such as occlusion, lighting variation, and dense crowd movement. Experimental results demonstrate that the system achieves approximately 95% counting accuracy in practical deployment scenarios.

Beyond real-time monitoring, the system incorporates a structured database architecture to support comprehensive historical analysis and predictive analytics. Occupancy data is aggregated at hourly, daily, weekly, and monthly levels, enabling peak-hour identification, trend analysis, and performance evaluation. Predictive analytics based on historical patterns and day-of-week analysis are implemented to forecast future occupancy levels and generate capacity alerts, transforming the system into an effective decision-support tool for library management.

The system is designed with a modular and scalable multi-camera architecture, allowing independent processing of multiple locations while providing a unified web-based dashboard for visualization. The dashboard presents live occupancy, historical trends, statistical summaries, and predictive insights through an intuitive interface accessible to non-technical users. By combining modern computer vision techniques, robust system architecture, and data analytics, this project demonstrates a practical and privacy-conscious solution for intelligent library occupancy monitoring and management.

Area of Study (Minimum 1 and Maximum 2): Computer Vision, Artificial Intelligence

Keywords (Minimum 5 and Maximum 10): People Counting, Occupancy Analytics, Computer Vision, YOLOv8, Object Tracking, Predictive Analytics, Library Management, Real-Time Monitoring

TABLE OF CONTENTS

TITLE PAGE	i
COPYRIGHT STATEMENT	ii
ACKNOWLEDGEMENTS	iii
ABSTRACT	iv
TABLE OF CONTENTS	v
LIST OF FIGURES	viii
LIST OF TABLES	xi
LIST OF SYMBOLS	xii
LIST OF ABBREVIATIONS	xi
CHAPTER 1 INTRODUCTION	1
1.1 Problem Statement and Motivation	2
1.2 Objectives	3
1.3 Project Scope and Direction	5
1.4 Contributions	9
1.5 Report Organization	10
1.6 Background Information	11
CHAPTER 2 LITERATURE REVIEW	13
2.1 Previous Work on People Detection & Tracking in Real-World	13
2.1.1 Wireless-Based People Detection and Counting	13
2.1.2 Hybrid and Sensor Fusion Approaches	15
2.1.3 Multi-Camera Fusion and Edge-Computing Architectures	16
2.2 Camera-Based Classical Computer Vision Methods	17
2.2.1 Introduction to Classical CV Approaches	17
2.2.2 Background Subtraction Techniques	17
2.2.3 Optical Flow Approaches	18
2.2.4 Contour and Blob Analysis	18
2.2.5 Comparative Evaluation of Classical CV Methods	19
2.3 Real-World Deployments in Libraries and Campuses	19

2.3.1 Thermal Imaging for Privacy-Preserving Occupancy Monitoring	20
2.3.2 Wi-Fi Sensing for Occupancy and Behavioral Insights	20
2.3.3 Vision-Based Seat Occupancy Systems	21
2.3.4 Hybrid Multi-Sensor Deployments	21
2.4 Transfer Learning with YOLOv8	22
2.4.1 Foundations of Transfer Learning in YOLOv8	22
2.4.2 YOLOv8 Network Architecture Overview	23
2.4.3 Applications in Remote Sensing	24
2.4.4 Applications in Wildlife Monitoring	25
2.4.5 Challenges and Future Directions	25
2.4.6 Summary	26
2.5 Critical Remarks of Previous Works	26
CHAPTER 3 SYSTEM METHODOLOGY/APPROACH (FOR DEVELOPMENT-BASED PROJECT)	28
CHAPTER 3 SYSTEM MODEL (FOR RESEARCH-BASED PROJECT)	
3.1 System Design Diagram/Equation	28
3.1.1 Core Counting Equations	28
3.1.2 Prediction Algorithms	29
3.1.3 Average Calculations	30
3.1.4 Percentage Change Calculations	30
3.1.5 Capacity Management	30
3.1.6 Object Tracking Equations	31
3.2 System Architecture Diagram	32
3.3 Use Case Diagram and Description	35
3.4 Activity Diagram	37
3.4.1 People Counting Process Flow	37
3.4.2 Dashboard Data Retrieval Activity Diagram	39
3.4.3 Daily Reset Process Flow	42
3.4.4 Predictive Analytics Calculation Flow	43

CHAPTER 4 SYSTEM DESIGN	46
4.1 System Block Diagram	46
4.1.1 Overall System Architecture Block Diagram	46
4.1.2 Data Flow Diagram	49
4.1.3 Software Architecture Layers	50
4.2 System Components Specifications	52
4.2.1 Hardware Components	52
4.2.1.1 IP Cameras	52
4.2.2 Software Components	53
4.2.2.1 Python Processing Server	53
4.2.2.2 PHP API Server	55
4.2.2.3 MySQL Database	57
4.2.2.4 Web Frontend	58
4.2.3 AI and Analytics Components	61
4.2.3.1 YOLOv8 Model	61
4.2.3.2 Object Tracking	62
4.2.3.3 Kalman Filter	62
4.2.3.4 YOLOv8 Training and Fine-Tuning Process	63
4.2.3.5 Average Analysis	71
4.2.3.6 Prediction Feature	72
4.3 Circuits and Components Design	74
4.3.1 Network Architecture	74
4.3.2 Data Flow Architecture	76
4.4 System Components Interaction Operations	77
4.4.1 Real-Time Processing Interaction	77
4.4.2 Database Interaction	79
4.4.3 Frontend-Backend Interaction	80
4.4.3.1 Real-Time Metrics Update	80
4.4.3.2 Video Streaming Interaction	81
4.4.3.3 Historical Data Retrieval	82
4.4.3.4 Prediction Calculation Interaction	83
4.4.4 Excel Export Interaction	84
4.4.5 Error Handling and Recovery	85

CHAPTER 5 SYSTEM IMPLEMENTATION (FOR DEVELOPMENT- BASED PROJECT)	87
CHAPTER 5 EXPERIMENT/SIMULATION (FOR RESEARCH- BASED PROJECT)	87
5.1 Hardware Setup	87
5.1.1 Processing Server / Computer	87
5.1.2 Network Infrastructure	88
5.2 Software Setup	88
5.2.1 Python Environment and Dependencies	88
5.2.2 CUDA and GPU Setup	89
5.2.3 YOLO/Ultralytics and Model Weights	89
5.2.4 Flask API Setup	89
5.2.5 PHP/Apache and MySQL Setup	89
5.2.6 Verification	90
5.3 Setting and Configuration	90
5.3.1 Camera and RTSP Configuration	90
5.3.2 Database Configuration	91
5.3.3 Application Parameters	91
5.3.4 Frontend and API Endpoints	91
5.3.5 Notes for Proper Configuration	92
5.4 System Operation (with Screenshot)	92
5.4.1 Startup Sequence	92
5.4.2 System Walkthrough	94
5.5 Implementation Issues and Challenges	100
5.5.1 Duplicate Counting and False Detection Issues	100
5.5.1.1 Duplicate Entry and Exit Counting	100
5.5.2 MJPEG Streaming Performance	101
5.5.3 Database Performance and Synchronization Issues	102
5.5.3.1 Database Write Performance Bottleneck	102
5.5.3.2 Daily Reset Synchronization	102
5.5.4 Real-Time Processing Performance Challenges	103
5.5.4.1 GPU Memory Management	103

5.5.4.2 Processing Latency and Frame Drops	103
5.5.5 Model Training and Fine-Tuning Challenges	104
5.5.5.1 Dataset Quality and Imbalance	104
5.5.5.2 Transfer Learning Hyperparameter Tuning	104
5.6 Concluding Remark	106
CHAPTER 6 SYSTEM EVALUATION AND DISCUSSION	107
6.1 System Testing and Performance Metrics	107
6.1.1 Testing Methodology Overview	107
6.1.2 Performance Metrics to Measure	108
6.1.2.1 Detection and Counting Accuracy Metrics	108
6.1.2.2 Predictive Analytics Accuracy	109
6.1.3 Testing Procedures	111
6.1.3.1 Accuracy Testing Procedure	111
6.1.3.2 Performance Testing Procedure	112
6.1.4 Metrics Summary Table	113
6.2 Testing Setup and Result	114
6.2.1 Testing Environment and Setup	114
6.2.2 Test Procedures	115
6.2.3 Test Results	116
6.2.3.1 Detection and Counting Accuracy	116
6.2.3.2 Performance Metrics	117
6.2.3.3 Predictive Analytics Results	117
6.2.4 Model Performance Comparison: Pre-trained vs Fine-tuned YOLOv8	118
6.2.4.1 Comparison Methodology	118
6.2.4.2 Detection Performance Comparison	122
6.3 Project Challenges	127
6.4 Objectives Evaluation	130
6.4.1 Evaluation Framework	130
6.4.2 Objective A: Technical Accuracy and Robustness	130
6.4.3 Objective B: Occupancy Analytics and Data Persistence	131

6.4.4 Objective C: Student Activity and Usage Pattern Analysis	131
6.4.5 Objective D: Usability and Dashboard Design	132
6.4.6 Objective E: AI Model Fine-Tuning and Integration	132
CHAPTER 7 CONCLUSION AND RECOMMENDATION	
7.1 Conclusion	136
7.2 Recommendation	137
REFERENCE	140
POSTER	145

LIST OF FIGURES

Figure Number	Title	Page
Figure 2.4.2	YOLOv8 Architecture	17
Figure 3.2	Architecture Diagram	19
Figure 3.3	Use Case Diagram	20
Figure 3.4.1	Activity Diagram for People Counting Process Flow	22
Figure 3.4.2	People Counting Process Flow Activity Diagram	23
Figure 3.4.3	Daily Reset Process Flow Activity Diagram	24
Figure 3.4.4	Predictive Analytics Calculation Flow Activity Diagram	25
Figure 4.1.1	System Block Diagram	19
Figure 4.1.2	Data Flow Diagram	20
Figure 4.1.3	Three-tier Software Architecture Diagram	21
Figure 4.2.2.4.1	Dashboard	35
Figure 4.2.3.4.1	Flow of the Fine-Tuning Process	38
Figure 4.2.3.4.2	Dataset Extracted from Recording	39
Figure 4.2.3.4.3	Labeled Data	40
Figure 4.2.3.4.4	Training Process	41
Figure 4.2.3.4.5	Best.pt Model Output	42
Figure 4.3.1.1	Network Architecture of the People Counter System	43
Figure 4.3.2.1	Data Flow Architecture of the Real-Time People Counter System	44
Figure 4.4.1.1	Sequence Diagram of Real-Time Processing Interaction	46
Figure 4.4.3.1.1	Real-Time Metrics Update Flow Between Frontend and Backend	48
Figure 4.4.3.2.1	Video Streaming Interaction Between Frontend and Backend	49
Figure 4.4.3.3.1	Historical Data Retrieval Interaction Between Frontend and Backend	50
Figure 4.4.3.4.1	Prediction Calculation Interaction Between Frontend and Backend	51
Figure 4.4.4.1	Excel Export Interaction Flow	52
Figure 5.4.1.1	Database in MySQL Service	54
Figure 5.4.1.2	Camera 1 Python Processing Server	54
Figure 5.4.1.3	Camera 2 Python Processing Server	55
Figure 5.4.1.4	XAMPP Control Panel	55
Figure 5.4.2.1	Dashboard View	56
Figure 5.4.2.2	History Counting View of Study Corner 1	56

Figure 5.4.2.3	Prediction Feature for Study Corner 1	57
Figure 5.4.2.4	Peak Hour Graph for Study Corner 1	58
Figure 5.4.2.5	Average Statistic View	58
Figure 5.4.2.6	Excel Reporting Export and View	59
Figure 5.4.2.7	Camera Selection in Average Feature	59
Figure 5.4.2.8	Camera Selection in Prediction Feature	60
Figure 6.1.2.1	Example of Annotated Ground-Truth Frame with Bounding Boxes	71
Figure 6.1.3.1	Example Test Cases Illustrating Occlusion and Multi-Person Scenarios	73
Figure 6.1.3.2	System Resource Usage and FPS Monitoring During Live Operation	74
Figure 6.2.2.1	Student Before Coming In	76
Figure 6.2.2.2	Two Students After Coming In	76
Figure 6.3	Line Chart Comparing Predicted and Actual Occupancy Trends	72

LIST OF TABLES

Table Number	Title	Page
Table 2.1	Specifications of laptop 4	4
Table 2.2	Comparison between the existing applications and proposed application	16
Table 3.1	Use Case Description of Register Account	18
Table 4.2	IP Camera Specifications and Functions	20
Table 4.2.2.1.1	Python Processing Server Specifications	30
Table 4.2.2.1.2	Server Configuration	30
Table 4.2.2.1.3	Python Server API Endpoints	31
Table 4.2.2.2.1	PHP API Server Specifications	32
Table 4.2.2.2.2	PHP API Files and Functions	32
Table 4.2.2.3.1	MySQL Database Specifications	33
Table 4.2.2.3.2	stats_log Table Structure	33
Table 4.2.2.3.3	daily_stats Table Structure	34
Table 4.2.2.4.2	Web Frontend Specifications	35
Table 4.2.2.4.3	Web Frontend Pages	35
Table 4.2.3.1.1	YOLOv8 Model Specifications	36
Table 4.2.3.1.2	Performance Characteristics	36
Table 4.2.3.2.1	Object Tracking Specifications	37
Table 4.2.3.3.1	Kalman Filter Specifications	37
Table 5.1.1	Processing Server/Computer Specifications and Functions	25
Table 5.1.2	Network Infrastructure	30
Table 6.1.2.1	Manual ground-truth counts versus system counts for selected test intervals on 14 December 2025	71
Table 6.2.4.2.1	Detection Performance Comparison Between Pre-trained YOLOv8x and Fine-tuned Model	77
Table 6.2.4.2.2	Detection Performance Across Different Operational Scenarios	78
Table 6.2.4.3.1	Entry Counting Accuracy Comparison	79
Table 6.2.4.3.2	Exit Counting Accuracy Comparison	80
Table 6.2.4.3.3	Overall System Counting Accuracy	81

LIST OF SYMBOLS

t	Time (Hour)
m	Number of samples/data points
i	Number of historical days / Index variable
k	Time step in Kalman filter
x	Position coordinate (x-axis)
y	Position coordinate (y-axis)
v_x	Velocity in x-direction
v_y	Velocity in y-direction
F	State transition matrix
P	Error covariance matrix
Q	Process noise covariance matrix
w	Process noise
K	Kalman gain
H	Observation matrix
R	Measurement noise covariance
z	Measurement vector
I	Identity matrix
Σ	Summation operator
\times	Multiplication operator
\geq	Greater than or equal to
$<$	Less than
\cdot	Dot product / multiplication

LIST OF ABBREVIATIONS

3D	Three-Dimensional
5G	Fifth Generation
AC	Alternating Current
AJAX	Asynchronous JavaScript and XML
AI	Artificial Intelligence
API	Application Programming Interface
ARIMA	AutoRegressive Integrated Moving Average
BLE	Bluetooth Low Energy
CNN	Convolutional Neural Network
COCO	Common Objects in Context
CPU	Central Processing Unit
CSS	Cascading Style Sheets
CSV	Comma-Separated Values
CUDA	Compute Unified Device Architecture
CV	Computer Vision
DHCP	Dynamic Host Configuration Protocol
DIOR	Dataset for Object Detection in Optical Remote Sensing
DOM	Document Object Model
ES6	ECMAScript 6
ESP8266	Espressif Systems 8266 (microcontroller)
FN	False Negative
FNR	False Negative Rate
FP	False Positive
FP16	16-bit Floating Point
FPN	Feature Pyramid Network
FPR	False Positive Rate
FPS	Frames Per Second
FYP	Final Year Project
FYP1	Final Year Project 1
FYP2	Final Year Project 2
GB	Gigabyte
GMM	Gaussian Mixture Model
GPIO	General Purpose Input Output
GPU	Graphics Processing Unit
GRU	Gated Recurrent Unit
HD	High Definition
HTML	HyperText Markup Language
HTML5	HyperText Markup Language Version 5

HVAC	Heating, Ventilation, and Air Conditioning
ID	Identifier
IoT	Internet of Things
IoU	Intersection over Union
IP	Internet Protocol
JSON	JavaScript Object Notation
KNN	K-Nearest Neighbors
LAN	Local Area Network
LSTM	Long Short-Term Memory
MAE	Mean Absolute Error
mAP	mean Average Precision
MB	Megabyte
MJPEG	Motion JPEG
ms	milliseconds
MySQL	My Structured Query Language
NMS	Non-Maximum Suppression
NVIDIA	NVIDIA Corporation (graphics processor manufacturer)
PAN	Path Aggregation Network
PDO	PHP Data Objects
PHP	Hypertext Preprocessor
PIR	Passive Infrared
PoE	Power over Ethernet
RAM	Random Access Memory
RFID	Radio-Frequency Identification
RGB	Red Green Blue
RL	Reinforcement Learning
RSSI	Received Signal Strength Indicator
RTSP	Real-Time Streaming Protocol
SMS	Short Message Service
SQL	Structured Query Language
SSD	Solid State Drive
UTAR	Universiti Tunku Abdul Rahman
VRAM	Video Random Access Memory
Wi-Fi	Wireless Fidelity
XLSX	Excel Open XML Spreadsheet
XML	Extensible Markup Language
YOLO	You Only Look Once
YOLOv8	You Only Look Once Version 8

Chapter 1: Introduction

1.0 Introduction

Modern libraries face increasing challenges in managing space utilization, capacity planning, and ensuring optimal service delivery. Traditional manual counting methods are labor-intensive, prone to human error, and provide only snapshot data that offers limited value for strategic decision-making. To address these limitations, this project presents an artificial intelligence-powered people counting and occupancy analytics system that automatically transforms live camera feeds into accurate, real-time entry and exit counts, alongside comprehensive historical insights and predictive analytics.

The system employs a sophisticated computer vision pipeline that integrates YOLOv8 object detection with ByteTrack multi-object tracking to achieve robust person identification and tracking across video frames. A carefully designed zone-transition algorithm ensures accurate counting by only registering entries and exits when individuals cross a virtual "gate" boundary, transitioning from outside to inside zones or vice versa. This approach incorporates anti-duplicate mechanisms including spatial suppression and temporal cooldown periods to minimize false counts, achieving approximately 95% accuracy in real-world conditions.

The technical architecture consists of three primary components: a Python-based processing engine that handles real-time video stream analysis with GPU acceleration, a Flask REST API backend that exposes live metrics and historical data endpoints, and a responsive web dashboard that provides intuitive visualization of current occupancy, historical trends, and predictive forecasts. A MySQL database serves as the persistent storage layer, maintaining both daily summary statistics and fine-grained time-series snapshots that enable comprehensive analytics including peak-hour identification, capacity alerts, and longitudinal pattern analysis.

Designed for deployment in library study corners or entrance corridors where people consistently pass-through well-defined passages, the system supports multi-camera configurations, allowing independent monitoring of multiple locations simultaneously. By seamlessly integrating real-time occupancy monitoring with historical trend analysis and predictive modeling, the system empowers library administrators to make data-driven decisions regarding crowd control, peak-time staffing allocation, capacity management, and long-term space planning. This project demonstrates how modern computer vision techniques

can transform raw video data into actionable, decision-ready intelligence for effective library management.

1.1 Problem Statement and Motivation

Libraries face persistent challenges in accurately counting occupants in real time, identifying peak usage patterns, and maintaining historical records for both operational planning and compliance. Manual methods such as clipboards, periodic checks, or estimates are labor-intensive, prone to human error, and do not scale well for large or dynamic environments. Studies have shown that manual counting methods can have error rates as high as 15-25% during peak hours, with accuracy degrading significantly when multiple staff members perform counts simultaneously or when counting intervals exceed 30 minutes. These inaccuracies compound over time, leading to unreliable historical data that undermines capacity planning and resource allocation decisions. The problem extends beyond simple counting accuracy. Traditional approaches fail to capture the temporal dynamics of library usage, missing critical patterns such as rapid occupancy spikes during class transitions, gradual build-ups during study periods, or unexpected crowd formations during special events. Without granular, time-stamped data, library administrators cannot optimize staffing schedules, predict peak demand periods, or implement effective crowd management strategies. This lack of visibility becomes particularly problematic in academic libraries.

While conventional surveillance systems can capture footage, they rarely convert it into structured, trustworthy metrics such as live occupancy, hourly distributions, or peak inside counts. Existing camera infrastructure typically serves security purposes rather than operational analytics, creating a missed opportunity to leverage existing investments for data-driven decision making. Furthermore, privacy concerns and data protection regulations require that any counting solution avoid storing identifiable personal information, making traditional access control systems or RFID-based tracking unsuitable for open-access library environments. The economic impact of inefficient occupancy management is substantial. Libraries that cannot accurately track usage patterns often over-staff during low-activity periods and under-staff during peak times, leading to both increased operational costs and degraded user experience. A study by the Association of Research Libraries found that libraries with poor occupancy visibility experienced 20-30% higher staffing costs per user interaction

compared to those with data-driven capacity management. Additionally, the inability to provide real-time occupancy information to users results in frustration when students arrive at a library only to find it at capacity, leading to reduced satisfaction and potential safety concerns during overcrowding.

This project is motivated by the need for an automated, accurate, and interpretable solution that can operate using existing camera infrastructure while addressing the specific challenges of library environments. By combining state-of-the-art computer vision models, YOLOv8 for detection and ByteTrack for tracking—with strict zone-based crossing logic, the system minimizes false counts and provides stable, near real-time insights. The motivation extends beyond efficiency, it seeks to enable data-driven decision-making by aligning staffing with demand, ensuring compliance with capacity limits, and optimizing library operations using objective, evidence-based analytics instead of guesswork.

The solution must also address the unique characteristics of library spaces, open-access policies that preclude physical barriers, diverse user behaviors ranging from quick book pickups to extended study sessions, and the need for non-intrusive monitoring that respects user privacy. Unlike retail or transportation environments where people move in predictable patterns, library users may linger, change direction, or engage in complex movement patterns that challenge traditional counting algorithms. The proposed system must therefore be robust enough to handle these real-world complexities while maintaining the accuracy and reliability necessary for operational decision-making.

1.2 Research Objectives

The primary objective of this project is to design and prototype a robust real-time people counting and occupancy analytics system tailored for library environments. To achieve this, four key goals are outlined.

First, the project emphasizes technical accuracy and robustness by developing a real-time people-counting pipeline that reliably detects and tracks individuals crossing a virtual gate. This involves integrating YOLOv8 with ByteTrack while incorporating safeguards to prevent double-counting and transient false detections. Second, the system will provide reliable occupancy analytics by consistently generating entry, exit, and current-inside metrics, persisting these in a relational database, and producing historical aggregations on hourly, daily,

weekly, and monthly bases. These features support peak-hour identification and long-term trend analysis.

Second, the project seeks to enable student activity tracking by extending the counting logic to analyze patterns of space utilization. Beyond simple entry and exit events, the system will capture temporal usage trends, such as peak study hours, average session duration, and day-to-day variations. These activity-based insights can inform library management on student behavior, resource allocation, and long-term planning.

Third, usability is prioritized through the development of a responsive web dashboard that presents live occupancy data, low-latency camera streams with MJPEG or JPEG fallback, and clear historical charts and tables. This ensures accessibility for non-technical users in operational settings. Finally, the project aims to ensure maintainability and resilience by implementing daily scheduled resets synchronized with opening hours, periodic database snapshots, GPU/CPU adaptability, and efficient memory management. These design choices enhance system stability during continuous operation in realistic conditions.

Fourth, the project emphasizes scalability and decision support. The system is designed to be modular and extensible, enabling future expansion to multi-camera fusion, predictive analytics, and cloud deployment without requiring major redesign. At the same time, it will serve as a decision-support tool for library management, translating raw occupancy and activity data into actionable insights for staffing optimization, space allocation, and long-term strategic planning.

Finally, AI Model Training and Adaptation, develop and fine-tune a machine learning model using datasets captured from the library environment to enhance detection accuracy and ensure adaptation to real-world operating conditions. This objective involves building a localized dataset of entry and exit events, transfer learning YOLOv8 (or an equivalent detection model) for domain-specific optimization and exploring predictive analytics techniques to forecast occupancy trends based on historical usage patterns. The dataset collection may begin in FYP1, while the actual training and adaptation process will be carried out in FYP2 as part of the extended system development.

1.3 Project Scope and Direction

Building upon the foundation established in FYP1, Final Year Project II (FYP2) extends the people counting system into a comprehensive, production-ready analytics platform. The scope encompasses multi-camera deployment, advanced predictive analytics, comprehensive historical reporting, and operational automation tools that transform the prototype into a practical solution for real-world library management.

- Multi-Camera Architecture and Scalability

FYP2 implements a scalable multi-camera architecture that supports simultaneous monitoring of multiple study corners or library entrances. Each camera operates as an independent processing instance with its own RTSP stream, dedicated MySQL database, and Flask API endpoint, enabling parallel operation without cross-interference. The system architecture supports Camera 1 (Study Corner 1) and Camera 2 (Study Corner 2) running on separate ports (8000 and 8001 respectively), with each maintaining its own `people_counter` and `people_counter_cam2` databases. The web dashboard provides unified visualization of both cameras through a responsive grid layout, allowing administrators to monitor multiple locations from a single interface. This modular design ensures that adding additional cameras requires minimal configuration changes—simply deploying a new camera script instance with a unique database identifier and port assignment.

- Predictive Analytics and Forecasting

A core enhancement in FYP2 is the implementation of predictive analytics capabilities that enable proactive decision-making. The system analyzes historical patterns from the last four weeks to generate multiple forecast types: next-hour occupancy predictions with confidence scores, today's remaining hours forecasts that project cumulative entries and exits for the rest of the day, and tomorrow's hourly predictions that estimate expected occupancy patterns for the following day. The predictive engine employs day-of-week pattern recognition, calculating averages for the same day of week across historical data to account for weekly cyclical patterns. Additionally, the system identifies peak hours for tomorrow, enabling administrators to anticipate high-traffic periods and allocate resources accordingly. Capacity alerts are generated

when predicted occupancy exceeds 80% of maximum capacity, providing early warning for potential overcrowding situations.

- Advanced Analytics and Statistical Reporting

FYP2 introduces comprehensive statistical analysis capabilities that transform raw counting data into actionable insights. The system calculates daily, weekly, and monthly averages across multiple metrics (entries, exits, peak inside counts), enabling trend identification and performance benchmarking. Peak daily records are tracked to identify exceptional traffic days, while day-of-week averages reveal weekly patterns that inform staffing and resource allocation decisions. Hourly averages for the current day of week provide granular insights into typical occupancy patterns throughout the day. The analytics engine also performs comparative analysis, comparing today's actual performance against historical averages and calculating percentage differences to highlight deviations from normal patterns. These statistics are exposed through dedicated API endpoints (`get_averages.php`) and visualized in the dashboard's "Average" view with interactive charts and comparison cards.

- Enhanced Historical Data Management

The historical data system has been significantly expanded to support multiple time granularities and efficient data retrieval. The system now supports hourly, daily, weekly, and monthly views, each optimized for different analytical purposes. Hourly views provide minute-by-minute granularity for detailed investigation, while daily, weekly, and monthly views aggregate data for trend analysis. Pagination has been implemented for large datasets, ensuring responsive performance even when querying months of historical data. The history pages (`history_counting.html` and `history_counting_cam2.html`) provide camera-specific views with interactive Chart.js visualizations, date selection controls, and period-based filtering. Data

is presented in both tabular and graphical formats, enabling administrators to choose the most appropriate representation for their analysis needs.

- Excel Export and Reporting

To support administrative reporting requirements, FYP2 includes a client-side Excel export feature that generates formatted monthly reports directly in the browser. The export functionality aggregates hourly entry data for a selected month and organizes it into time-block categories (8am-1pm, 1pm-5pm, 5pm-7pm, 7pm-8pm, 8pm-9pm, 9pm-10pm, 10pm-11pm) for each day, providing a structured format suitable for library management reports. The system uses the SheetJS (xlsx.js) library to generate XLSX files without requiring server-side processing, ensuring privacy and reducing server load. The export interface includes a preview table that displays the formatted data before download, allowing users to verify the content and structure. Grand totals are automatically calculated and included in the exported file, providing summary statistics for the entire month.

- Automated System Management

To support automated system startup and reduce manual intervention, a Windows batch script was developed to initialize all required services and applications for the FYP system. When executed, the script first starts the Apache web server in minimized mode, followed by the MySQL database service, ensuring that the backend web and database components are available before other processes begin. Short delays are introduced between service startups to allow each service sufficient time to initialize properly. After the web and database services are running, the script launches two Python-based camera processing programs located on the D drive, which are responsible for real-time video analysis and people counting from multiple camera inputs. These Python scripts are executed in separate command windows to allow continuous monitoring and logging during runtime. Finally, the script automatically opens the administrator dashboard through a local HTML file hosted on the Apache server, providing immediate access to system monitoring and control features. This automated startup

mechanism enables the entire system to be launched with a single action, improving operational efficiency, supporting system recovery, and simplifying deployment for non-technical staff.

- Model Training and Customization Tools

To support model optimization and adaptation to specific environments, FYP2 includes a comprehensive suite of transfer learning and dataset preparation tools. The `train_yolo.py` script enables fine-tuning of pre-trained YOLO models on custom datasets, supporting configurable training parameters including epochs, batch size, learning rate, and layer freezing strategies. Dataset organization utilities (`organize_dataset.py`) facilitate the creation of train/validation splits from extracted video frames, with optional intelligent selection that balances positive and negative examples. The `auto_label_frames.py` script accelerates the labeling process by generating initial YOLO-format annotations using a pre-trained model, which can then be manually refined. Frame extraction tools (`extract_training_frames.py`) enable systematic dataset creation from recorded video, supporting both interval-based and count-based extraction strategies. These tools collectively enable the system to be adapted to new environments and improve accuracy through domain-specific training.

- Enhanced Dashboard and User Experience

The dashboard has been redesigned to support multi-camera monitoring with improved navigation and visualization capabilities. The main dashboard (`admin_dashboard.html`) features a camera grid layout that displays real-time metrics and live video feeds for both cameras simultaneously. Navigation has been enhanced with dedicated views for "Head Counting Current", "History Counting", "Peak Hours Graph", "Prediction", and "Average" analytics. Each view is camera-aware, with dropdown selectors allowing users to switch between Camera 1 and Camera 2 data. The prediction view displays next-hour forecasts, today's remaining predictions, tomorrow's hourly forecasts, peak hour identification, and capacity alerts in an intuitive card-based layout. The average view presents comprehensive

statistical summaries with comparison metrics and trend indicators. All views maintain responsive design principles and provide real-time updates through AJAX polling.

- Database Architecture and Data Persistence

The database architecture has been extended to support multi-camera deployments while maintaining data isolation and query performance. Each camera maintains its own database instance (`people_counter` and `people_counter_cam2`), ensuring that data from different locations remains separate and queries remain efficient. The database schema includes `stats_log` tables for per-minute time-series data and `daily_stats` tables for aggregated daily summaries. Indexing strategies have been optimized for common query patterns, particularly for time-based filtering and day-of-week calculations used in predictive analytics. The system implements automatic daily reset functionality that clears counters at midnight while preserving historical data, ensuring accurate daily statistics without manual intervention.

1.4 Contributions

This project makes significant contributions by bridging advanced computer vision techniques with practical library facility management through a production-ready automated occupancy analytics system. It integrates YOLOv8 object detection with ByteTrack multi-object tracking to achieve approximately 95% counting accuracy in real-world conditions, operating in real time at 25 - 30 FPS. A strict zone-transition counting strategy (outside → gate → inside), combined with Kalman filtering, appearance-based histogram matching, and multiple anti-duplicate safeguards, ensures reliable entry and exit detection despite occlusions, lighting variations, and dense crowd scenarios. Beyond real-time monitoring, the system establishes a scalable data architecture that persistently stores fine-grained time-series logs and aggregated statistics, enabling average occupancy analysis, peak-hour identification, and long-term trend evaluation across hourly, daily, weekly, and monthly intervals. These historical datasets form the basis for predictive analytics, supporting future extensions such as occupancy forecasting and demand estimation to assist proactive decision-making. The system is designed for continuous 24/7 operation with automated recovery, GPU/CPU adaptability, scheduled resets, and optimized memory management, ensuring robustness and maintainability. A responsive

web dashboard further enhances usability by presenting live occupancy, historical trends, and visual analytics in an intuitive format accessible to non-technical users. Overall, the project delivers a modular, extensible foundation for intelligent occupancy monitoring, transforming raw video streams into actionable insights that support operational efficiency, resource allocation, and long-term strategic planning in library environments.

1.5 Report Organization

This report is organized into seven chapters to present the development and evaluation of the proposed people counting and occupancy analytics system in a clear and systematic manner. Chapter 1 introduces the project by outlining the background, problem statement, research objectives, project scope, contributions, and motivation. It provides the context for applying computer vision–based people counting in library environments and establishes the significance of the study.

Chapter 2 presents a comprehensive literature review of existing approaches related to people detection, tracking, and occupancy monitoring. It examines classical computer vision techniques, wireless and sensor-based methods, deep learning–based solutions, and real-world deployments in libraries and campuses. This chapter highlights the strengths and limitations of previous works and justifies the adoption of YOLOv8-based computer vision techniques for this project.

Chapter 3 describes the system methodology and overall approach. It details the system architecture, mathematical models, counting logic, prediction algorithms, and object tracking mechanisms. Use case diagrams, activity diagrams, and architectural explanations are provided to illustrate how the system operates from video input to data visualization and analytics.

Chapter 4 focuses on system design, presenting the block diagrams, data flow, software architecture layers, and component interactions. This chapter explains how the hardware, software, databases, and APIs are integrated into a scalable multi-camera system capable of real-time processing and historical data management.

Chapter 5 discusses the system implementation in detail. It covers hardware and software setup, configuration procedures, system operation, implementation challenges, and practical

issues encountered during development. Screenshots and operational explanations are included to demonstrate the functionality of the implemented system.

Chapter 6 presents the system evaluation and discussion. It describes the testing setup, performance metrics, experimental results, and analysis of system accuracy and reliability. The chapter also evaluates the project objectives and discusses challenges faced during development and deployment.

Finally, Chapter 7 concludes the report by summarizing the overall project outcomes and key contributions. It also provides recommendations for future improvements and enhancements, highlighting potential directions for extending the system's accuracy, scalability, and analytical capabilities in real-world library environments.

1.6 Background Information

The rapid growth of digital technologies has enabled organizations and institutions to move beyond manual processes and adopt automated solutions to improve efficiency. One area that has benefited greatly from such advances is people counting and occupancy monitoring. Traditionally, libraries and public facilities relied on manual headcounts or mechanical turnstiles to estimate the number of visitors. While simple, these methods are often inaccurate, labor-intensive, and incapable of providing real-time insights. With the rise of computer vision and artificial intelligence, automated systems now offer more precise, scalable, and non-intrusive ways to measure occupancy and understand usage patterns.

Computer vision, a subfield of artificial intelligence, focuses on enabling machines to interpret and understand visual information from the environment. Early techniques relied on handcrafted methods such as background subtraction, contour detection, and optical flow to identify people in video streams. While useful, these methods struggled in complex environments where lighting, occlusion, and crowd density varied. The advent of deep learning, particularly convolutional neural networks (CNNs), revolutionized the field by significantly improving object detection and tracking accuracy. Modern approaches such as YOLO (You Only Look Once) and multi-object tracking algorithms can now deliver real-time,

high-accuracy results even in crowded or dynamic spaces, making them ideal for library environments.

Within libraries, occupancy monitoring plays an important role in resource allocation, safety, and user satisfaction. Libraries often serve as study hubs, event spaces, and community centres where foot traffic varies throughout the day. Accurate people counting allows administrators to identify peak usage hours, ensure compliance with safety regulations, and allocate staff more effectively. At the same time, aggregated occupancy data helps with long-term planning, such as deciding on seating arrangements, study room availability, or infrastructure expansion. Importantly, modern systems must also respect user privacy by avoiding the storage of personally identifiable information, focusing instead on anonymous metrics such as entry counts and occupancy levels.

Key technical terms relevant to this project include detection tracking, and occupancy analytics, aggregating counts over time into meaningful metrics such as current inside, entries, exits, and peak hours. Understanding these concepts provides the foundation for following the later chapters of this report, where the integration of computer vision, databases, and web technologies is described in detail. By combining insights from artificial intelligence, software engineering, and data analytics, this project seeks to demonstrate how modern computer vision can be applied to address practical challenges in library management.

CHAPTER 2

Literature Reviews

People counting and occupancy monitoring have become increasingly important in public and semi-public spaces [1] such as libraries, airports, shopping malls, and universities. Accurate knowledge of how spaces are being used enables administrators to make informed decisions about capacity management, staffing, safety compliance, and resource allocation. Traditionally, occupancy has been estimated through manual headcounts or basic access logs [2], but these methods are labor-intensive, error-prone, and lack the temporal resolution needed for modern operational planning.

In recent years, a wide range of technological solutions have been explored to overcome these limitations. Sensor-based methods, including infrared beams, RFID tags, and Wi-Fi/Bluetooth tracking [3], offer partial automation but suffer from accuracy, cost, or privacy concerns. Camera-based approaches, which leverage classical computer vision techniques such as background subtraction and optical flow, have provided more direct visibility but often struggle with robustness under real-world conditions like occlusion, lighting variations, and crowded environments. [2]

This chapter reviews the existing literature and practices in people counting and occupancy analytics. It begins with an examination of traditional manual methods and sensor-based approaches, followed by camera-based techniques, including both classical computer vision and deep learning models [3]. The discussion then considers the role of database storage and dashboard systems in translating raw data into decision-ready insights [1]. Finally, the chapter provides a critical analysis of these approaches, identifying their strengths, weaknesses, and gaps, and explains how the proposed system builds on these foundations to deliver an AI-driven, real-time occupancy analytics solution for library environments.

2.1 Previous Work on People Detection & Tracking in Real-World

The problem of accurately detecting and tracking people in real-world environments has attracted attention from researchers, industry, and public institutions for decades [4]. The challenge stems from diverse application domains, ranging from public safety and transportation hubs to libraries, classrooms, and retail spaces. Prior works can be broadly grouped into wireless-based approaches, hybrid sensor fusion systems, and multi-camera or edge-computing frameworks. Each family of approaches brings unique advantages and limitations depending on the scale, infrastructure, and performance requirements of the target environment [5].

2.1.1 Wireless-Based People Detection and Counting

Wireless methods have gained popularity due to their low infrastructure cost and ability to leverage existing Wi-Fi and Bluetooth devices. They typically rely on signal fluctuations or device presence to infer occupancy levels.

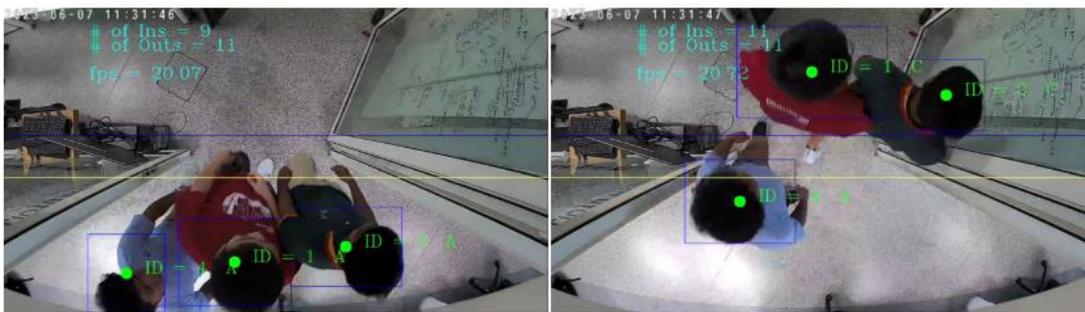


Figure 2.1 Wireless-Based People Detection and Counting Diagram

One prominent example is AFOROS by Gómez et al. [19], a Wi-Fi probe-based system designed for crowd monitoring in public spaces. AFOROS achieved up to 95% accuracy by capturing the number of Wi-Fi-enabled devices in an environment, making it suitable for large events such as concerts or conferences. Figure 2.1, example of a real-time deployment result from . Before entering the room, the count is 9. After two individuals (one in a red T-shirt and one in a green T-shirt) enter, the “# of Ins” value increases by 2, confirming their entry. Their IDs also change from A to C, reflecting the transition from outside to inside. However, the system assumes that most users carry Wi-Fi-enabled devices, which is not always true in diverse settings such as libraries, where some users may disable connectivity or use devices without Wi-Fi.

Similarly, Liu et al. [20] introduced CrossCount, which combines temporal patterns in Wi-Fi links with deep learning models to distinguish between transient and persistent users.

Their system showed strong performance with an error margin of less than two individuals in classroom-scale deployments. Importantly, it demonstrated how wireless-based systems can move beyond simple device counts toward more refined modelling of human presence.

In the Bluetooth domain, Hossain et al. [21] evaluated BLE beacon-based occupancy estimation, finding that it could maintain an error margin of fewer than 0.5 people on average in controlled indoor trials. BLE offers lower energy consumption compared to Wi-Fi and can be deployed in beacon-based infrastructures. However, its reliance on people carrying compatible devices again limits universality.

Other researchers have explored Radio Signal Strength Indicator (RSSI) analysis. Zhang et al. [22] used RSSI fluctuations with machine learning models to estimate occupancy in meeting rooms, achieving more than 98% accuracy for up to nine individuals. This illustrates the potential of wireless signals as a proxy for presence, though scalability and robustness remain key challenges.

Wireless methods excel in cost-effectiveness, scalability, and non-intrusiveness, but they are often criticized for their dependence on device availability, potential privacy issues (device tracking), and reduced accuracy in highly dynamic or sparse environments.

2.1.2 Hybrid and Sensor Fusion Approaches

Hybrid systems attempt to overcome the weaknesses of individual sensing modalities by combining visual, environmental, and wireless data.

For example, Mihailescu et al.[5] proposed a multimodal classroom occupancy detection system that integrated overhead video analysis with gas sensors monitoring CO₂ levels. By correlating environmental changes with visual data, the system achieved high accuracy even under conditions where occlusion or lighting variation would normally hinder purely vision-based models.

Erickson et al. [23] demonstrated another form of hybrid sensing by using depth cameras alongside environmental data (temperature, CO₂, and humidity) in smart buildings. Their long-term experiments revealed that occupancy detection could improve HVAC efficiency while maintaining comfort and reducing energy costs by up to 18%.

Hybrid approaches are also common in transportation hubs. For example, airports often use infrared sensors at gates combined with CCTV systems to improve people flow monitoring. A notable project is London Heathrow's Passenger Flow Management System, which integrates multiple sensing modalities to minimize bottlenecks during peak hours[24].

The strength of hybrid systems lies in robustness and redundancy—when one modality underperforms (e.g., cameras at night), others can compensate. However, this robustness comes at a cost: hybrid systems often require substantial hardware investment, integration expertise, and careful calibration, making them less practical for environments with limited budgets such as academic libraries.

2.1.3 Multi-Camera Fusion and Edge-Computing Architectures

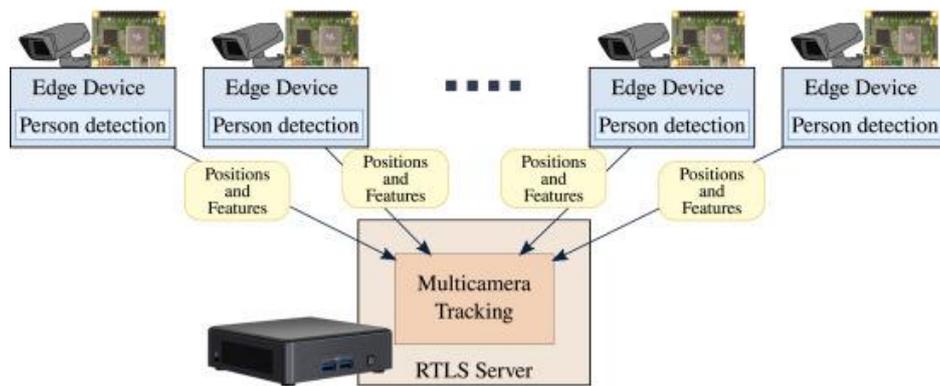


Figure 2.2 Multicamera edge-computing system

In crowded or complex environments, single-sensor or single-camera solutions are insufficient. Multi-camera and edge-computing frameworks are therefore widely explored.

Guo et al. [25] introduced Zensors++, a system that employs overlapping video feeds combined with density estimation algorithms to produce accurate real-time crowd counts. Unlike device-based methods, Zensors++ directly observes individuals, making it effective in environments where not all users carry devices. However, the system required extensive computational resources and raised privacy concerns since video feeds were centrally processed.

To address privacy and scalability, Jin et al. [26] designed a fog-computing occupancy monitoring system for university libraries. Their system processed video data at the edge, extracting only anonymized features (e.g., bounding boxes, counts) before transmitting results to the central server. This reduced both bandwidth and privacy risks while maintaining low-latency reporting. The project demonstrated a realistic balance between surveillance needs and ethical considerations in sensitive environments like libraries.

From an industrial perspective, Cisco’s Smart Campus initiative [27] integrates multi-camera fusion with IoT devices to monitor occupancy across university facilities. Their system emphasizes compliance (ensuring occupancy does not exceed fire safety limits), resource

allocation (adjusting HVAC/lighting), and operational efficiency. While effective, it highlights the high infrastructure investment typical of enterprise-grade solutions.

Multi-camera and edge-computing solutions are the most scalable and accurate but require substantial investment in hardware, networking, and processing capacity. Additionally, ethical challenges regarding data protection remain unresolved in many deployments.

2.2 Camera-Based Classical Computer Vision Methods

2.2.1 Introduction to Classical CV Approaches

Before the rise of deep learning, people detection and counting in video surveillance primarily relied on classical computer vision (CV) techniques. These methods used pixel-level motion cues, such as background subtraction and optical flow, or handcrafted features, such as edges, silhouettes, and contours, to identify moving individuals in video streams. Their main advantage was low computational cost and real-time feasibility on standard CPUs, making them suitable for deployment before the widespread availability of GPUs. Despite their efficiency, these methods often struggled with robustness when applied to real-world environments, where variable lighting, occlusion, and crowd density significantly degraded their performance.

2.2.2 Background Subtraction Techniques

Background subtraction is one of the earliest and most commonly applied approaches in classical CV. The technique involves building a model of the static scene and identifying moving objects as foreground whenever deviations occur. Early implementations relied on simple frame differencing, where two consecutive frames were subtracted to highlight motion. While computationally efficient, this method produced noisy results and suffered from false positives caused by camera jitter or illumination variations.

A major advancement came with the introduction of Gaussian Mixture Models (GMM) by Stauffer and Grimson [6], which modelled each pixel as a mixture of Gaussian distributions. This innovation allowed the background model to adapt to dynamic environments with repetitive motions such as waving trees or flickering monitors and became the foundation for

OpenCV's. Zivkovic[9] further refined background subtraction by introducing a K-Nearest Neighbours (KNN) model that employed non-parametric pixel histories to capture background dynamics. To address the common issue of shadows being misclassified as moving objects, Horprasert et al [7] proposed using chromaticity and brightness cues for more reliable shadow suppression.

In practice, background subtraction has been applied to airport passenger flow monitoring and traffic surveillance systems, where environmental conditions are relatively stable. For example, Benabbas et al.[11] demonstrated the use of background subtraction in a controlled pedestrian zone to perform crowd counting. However, the limitations of this approach remain notable. It is highly sensitive to sudden lighting changes, struggles with dynamic backgrounds such as reflections, and fails to properly track stationary individuals who blend into the background over time.

2.2.3 Optical Flow Approaches

Another widely adopted classical CV technique is optical flow, which estimates the apparent motion of pixels across consecutive frames to produce a vector field representing both magnitude and direction of movement.

Building on these foundations, Benabbas et al. [11] proposed a spatial-temporal optical flow method that constructed motion history images across a virtual line, allowing the system to estimate the number of individuals crossing a scene. Cheng et al. [10] further advanced the technique by applying density-based clustering to optical flow vectors, enabling the segmentation of crowd movements into groups and supporting higher-level behavioral analysis.

The strengths of optical flow methods include their independence from object appearance, making them suitable for low-resolution or grayscale video, and their ability to capture fine-grained motion. However, they come with several weaknesses. Optical flow cannot effectively distinguish overlapping individuals in dense crowds because motion vectors tend to merge. The method is also computationally more demanding than background subtraction and is vulnerable to noise caused by camera vibration or background motion.

2.2.4 Contour and Blob Analysis

Contour and blob analysis represent another line of research in classical CV. These methods typically begin with background subtraction to obtain a foreground mask, after which contours or blobs of moving objects are extracted. By analysing the size, shape, and aspect ratio of these blobs, systems could classify them as human figures. A well-known example is the W4 system developed by Haritaoglu et al. [8], which tracked people and their activities by combining silhouette-based features with contour analysis.

In practice, contour and blob methods were applied to early surveillance systems in subway stations and retail environments, where blob size was directly correlated with the number of individuals passing through a monitored entrance. However, these methods faced significant limitations. In dense environments, blobs often merged into a single region, causing undercounting. Furthermore, they were unable to maintain identity across frames, making them unsuitable for applications where direction-aware entry and exit tracking was required. Occlusion also posed a major challenge, as overlapping individuals could not be separated reliably.

2.2.5 Comparative Evaluation of Classical CV Methods

Several studies have compared the effectiveness of background subtraction, optical flow, and contour analysis. In general, background subtraction tends to perform best in stable environments with controlled lighting and static backgrounds, while optical flow provides superior continuity of motion in dynamic crowd scenes. Contour and blob analysis, although conceptually simple, fails under heavy occlusion and dense crowding.

For instance, Zhang et al. [12] evaluated GMM-based background subtraction and optical flow approaches on the PETS2009 pedestrian dataset. Their findings revealed that background subtraction achieved higher accuracy when detecting moving individuals frame by frame, while optical flow offered more reliable tracking of continuous motion. These comparative evaluations underscore the strengths and weaknesses of classical CV methods and highlight why they were eventually surpassed by deep learning techniques.

2.3 Real-World Deployments in Libraries and Campuses

While much of the research on people detection and occupancy monitoring has been carried out in controlled laboratory conditions, real-world deployments in libraries and campuses present unique challenges. These include privacy concerns, high-density crowds, variable lighting, and the need for scalable and cost-effective solutions. This section reviews

actual implementations of such systems in higher education environments, focusing on thermal imaging, Wi-Fi sensing, vision-based solutions, and hybrid multi-sensor approaches.

2.3.1 Thermal Imaging for Privacy-Preserving Occupancy Monitoring

Thermal imaging sensors have emerged as a promising solution for occupancy monitoring in sensitive spaces like libraries, where privacy is a major concern. Unlike RGB cameras, thermal sensors capture heat signatures, thus avoiding personally identifiable imagery. Wang and Shao [13] conducted a longitudinal deployment of thermal sensors in a UK university library, monitoring occupancy continuously over an academic year. Their study revealed discrepancies between perceived and actual space utilization, such as underuse during weekends and unexpected activity during vacation periods. The authors demonstrated that thermal imaging is highly effective in detecting occupancy trends without compromising user anonymity.

This work highlights the suitability of thermal imaging for long-term studies of user behaviour in academic buildings. However, the approach is limited by higher hardware costs compared to Wi-Fi-based solutions, and its effectiveness decreases in open spaces where heat signatures overlap. Nevertheless, it serves as an important example of balancing privacy preservation with the need for granular, reliable occupancy data in real-world library environments.

2.3.2 Wi-Fi Sensing for Occupancy and Behavioural Insights

Wi-Fi infrastructure has been repurposed in multiple universities to track student occupancy and movement behaviour. By leveraging probe request signals from mobile devices, these systems infer how many users are present and how they move across different spaces. Wang and Shao [14] conducted a 30-day monitoring study in a university library using Wi-Fi indoor positioning. Their analysis revealed distinct occupancy patterns, including correlations between arrival times and study durations, as well as variations across weekdays and weekends. Such insights provide administrators with data to adjust opening hours, allocate resources, and optimize energy use.

Galluzzi et al. [15] further demonstrated a low-cost Wi-Fi sniffer-based occupancy estimation system deployed across multiple campus buildings. Using ESP8266 microcontrollers, the system collected probe request signals and applied supervised machine learning models to classify occupancy states in real time. Results showed high accuracy at a

fraction of the cost of traditional sensing infrastructures, making it a scalable option for libraries with limited budgets.

These Wi-Fi-based deployments illustrate how existing network infrastructure can be exploited for large-scale behavioural analytics. However, they are inherently dependent on device-carrying behaviour.

2.3.3 Vision-Based Seat Occupancy Systems

Camera-based solutions remain a widely studied approach due to their ability to provide fine-grained spatial information, such as seat-level occupancy. Yang et al. [16] proposed a dual-channel Faster R-CNN model for library seat detection, which combines virtual data with real images to overcome dataset limitations. Their system was deployed in a university library and integrated into a web and mobile dashboard to provide real-time seat availability updates to students. The evaluation showed that the system improved recognition accuracy while reducing computational cost, thereby enhancing the usability of vision-based approaches in dense, real-world library settings.

Compared to Wi-Fi or thermal imaging, camera-based systems offer higher spatial resolution and seat-level precision, enabling real-time visualization of space usage. However, they face challenges related to privacy concerns, computational demands, and occlusion handling in crowded environments. Despite these drawbacks, such systems are increasingly deployed in smart libraries to support seat reservation, space management, and user satisfaction.

2.3.4 Hybrid Multi-Sensor Deployments

To address the limitations of single-sensor systems, researchers have explored hybrid approaches that combine multiple sensing modalities for more robust and scalable occupancy monitoring. Kim et al. [17] deployed a multi-sensor fusion system in a university campus, combining passive infrared (PIR) sensors, Wi-Fi signals, and overhead cameras to estimate real-time occupancy. Their findings showed that multi-sensor fusion significantly outperformed single-sensor methods, reducing false positives from shadows in vision systems and compensating for the biases of Wi-Fi-based tracking.

Similarly, Yang et al. [18] presented a hybrid IoT-based occupancy framework integrating environmental sensors with computer vision to predict occupancy levels in

academic study rooms. The combination of environmental and vision data provided higher robustness in low-light or occluded scenarios, achieving more consistent results across different room layouts and times of day.

Hybrid systems demonstrate the potential of sensor complementarity, where the weaknesses of one modality are mitigated by the strengths of another. However, they also introduce challenges such as increased infrastructure costs, sensor calibration, and data integration complexity. Nonetheless, these systems represent a promising direction for libraries and campuses aiming for high-accuracy, privacy-compliant, and scalable monitoring solutions.

2.3.5 Summary of Real-World Deployments

The reviewed case studies illustrate how libraries and campuses are adopting heterogeneous sensing strategies depending on their constraints and priorities. Thermal imaging [13] ensures privacy while capturing long-term occupancy trends, Wi-Fi sensing [14], [15] leverages existing infrastructure for cost-effective monitoring, vision-based systems [16] provide seat-level precision but raise privacy and scalability concerns, and hybrid multi-sensor deployments [17], [18] deliver robust and reliable performance through sensor fusion. Together, these deployments highlight the trade-offs between accuracy, privacy, cost, and scalability, shaping the evolution of smart campus and library management systems.

2.4 Transfer Learning with YOLOv8

Transfer learning has become a widely adopted strategy in deep learning, particularly for computer vision tasks where collecting large, labeled datasets is costly or impractical. By leveraging pretrained models, transfer learning enables knowledge gained from large-scale datasets to be reused and adapted to new application domains. YOLOv8, as the latest generation in the "You Only Look Once" object detection family, offers strong pretrained weights and flexible fine-tuning mechanisms, making it well suited for transfer learning. This section reviews the principles of transfer learning in YOLOv8 and examines its applications across different domains, along with key challenges and future research directions.

2.4.1 Foundations of Transfer Learning in YOLOv8

YOLOv8 is typically pretrained on large-scale benchmark datasets such as MS COCO, allowing the model to learn generalized visual features including edges, textures, and basic

object structures. In transfer learning, these pretrained features are reused as a foundation for domain-specific tasks. A common practice involves freezing the early layers of the backbone network to retain low-level feature extraction while fine-tuning the later layers and detection heads to adapt to the target domain[33]. This approach helps stabilize training, reduces the risk of catastrophic forgetting, and improves convergence when training data is limited.

In addition to layer freezing strategies, researchers have highlighted the importance of input resolution and model configuration when performing transfer learning with YOLOv8[30]. Higher input resolutions can improve detection performance, particularly for small objects, while domain-specific data augmentation techniques such as rotation, mosaic augmentation, and brightness adjustment help reduce the gap between source and target domains[31]. These strategies collectively enhance YOLOv8’s adaptability and robustness in transfer learning scenarios.

2.4.2 YOLOv8 Network Architecture Overview

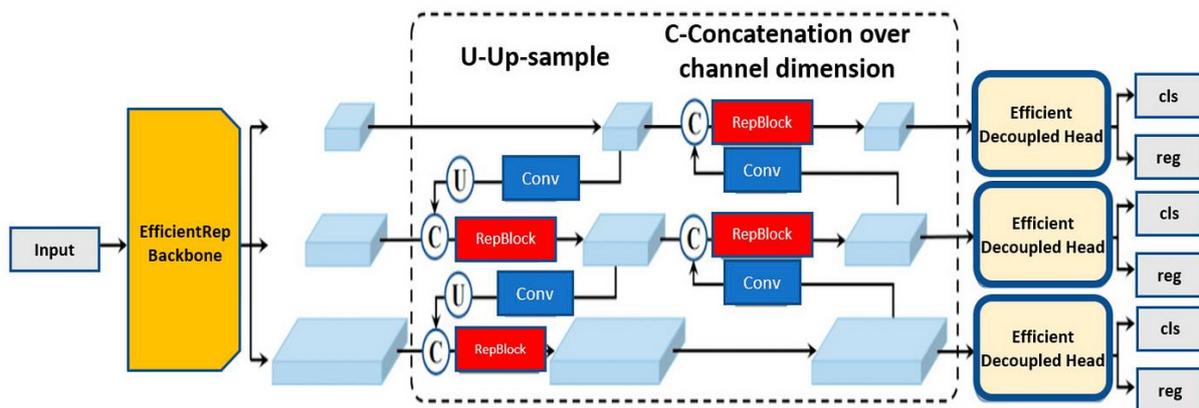


Figure 2.4.2 Yolov8 architecture[34]

Figure 2.4.2 illustrates the overall architecture of the YOLOv8 object detection framework, which is designed to balance detection accuracy and real-time efficiency. The architecture follows a three-part structure consisting of an EfficientRep backbone, a feature aggregation neck, and an efficient decoupled detection head. This design enables YOLOv8 to achieve strong performance across multiple object scales while maintaining low latency, making it suitable for real-time applications such as crowd monitoring and occupancy analytics.

The EfficientRep backbone is responsible for extracting hierarchical feature representations from the input image. It employs a series of convolutional layers and RepBlocks, which are

optimized structural blocks that combine multi-branch convolutions during training and re-parameterize into a single-path structure during inference. This approach enhances feature extraction capability while reducing inference-time computational cost. As a result, the backbone efficiently captures both low-level spatial details and high-level semantic information necessary for robust object detection.

Following feature extraction, YOLOv8 utilizes a feature aggregation neck that adopts an up-sampling and concatenation strategy inspired by Feature Pyramid Networks (FPN) and Path Aggregation Networks (PAN). The up-sampling operations (U) increase the spatial resolution of deeper feature maps, while concatenation (C) along the channel dimension fuses features from different network depths. This multi-scale feature fusion allows the model to better detect objects of varying sizes, particularly small and medium-sized targets that may otherwise be missed in deeper layers. RepBlocks and convolution layers within the neck further refine the fused features and enhance contextual understanding.

At the final stage, YOLOv8 employs an efficient decoupled head, which separates classification (cls) and regression (reg) tasks into independent branches. Unlike earlier YOLO versions that used a coupled head, this decoupled design reduces task interference and improves convergence stability during training. The regression branch focuses on precise bounding box localization, while the classification branch predicts object categories independently. This separation has been shown to improve both localization accuracy and classification confidence, especially in crowded or complex scenes.

2.4.3 Applications in Remote Sensing

Remote sensing imagery presents distinct challenges due to its top-down viewing angle, dense object distributions, and the small size of targets within large scenes. Transfer learning with YOLOv8 has been widely applied to remote sensing datasets such as DIOR and various ship detection benchmarks. Empirical studies demonstrate that fine-tuning pretrained YOLOv8 models significantly outperforms training from scratch, particularly in terms of recall and precision for small objects such as ships, vehicles, and aircraft[31] [32].

Researchers have shown that increasing input resolution and applying augmentations that mimic aerial imaging conditions, including rotations and perspective transformations, effectively reduce domain shift. Moreover, transfer learning substantially decreases training time and data requirements, making YOLOv8 a practical solution for remote sensing applications where annotated data is scarce and computational resources may be limited.

2.4.4 Applications in Wildlife Monitoring

Wildlife monitoring is another domain that benefits significantly from transfer learning due to limited labeled datasets and severe class imbalance. YOLOv8 has been adapted for wildlife detection and species monitoring tasks, where pretrained weights enable stable and efficient training on small datasets. Studies indicate that transfer learning allows YOLOv8 to achieve competitive performance compared to traditional convolutional neural network classifiers such as ResNet and DenseNet.[28] [29]

Beyond detection, YOLOv8 also supports detection–classification pipelines in which bounding boxes generated by the detector are further processed by specialized classifiers for fine-grained species recognition[28]. Strong data augmentation strategies play a crucial role in these applications, helping to mitigate overfitting and improve generalization across varying environmental conditions, such as changes in lighting, background, and animal posture.

2.4.5 Challenges and Future Directions

Despite the demonstrated effectiveness of transfer learning with YOLOv8, several challenges remain. One major issue is domain shift, where significant visual differences between the pretraining dataset (e.g., COCO) and the target domain limit transferability. Small-object detection continues to be difficult when objects occupy only a few pixels, often requiring multi-scale feature extraction strategies or specialized architectural modifications.

Additionally, limited labeled data remains a persistent challenge in many application domains. Emerging approaches such as semi-supervised learning, active learning, and synthetic data generation have been proposed as complementary techniques to enhance transfer learning performance. Future research directions also include multi-modal integration, such as

combining RGB imagery with thermal or multispectral data, to improve robustness under challenging environmental conditions.

2.4.6 Summary

In summary, transfer learning with YOLOv8 has demonstrated substantial benefits across a wide range of application domains, particularly in scenarios with limited labeled data. By reusing pretrained features, selectively fine-tuning network components, and applying domain-specific augmentation strategies, researchers have achieved improvements in detection accuracy, training efficiency, and model generalization. Continued advancements in cross-domain pretraining, semi-supervised learning, and multi-modal data fusion are expected to further enhance YOLOv8's adaptability and effectiveness in real-world deployments.

2.5 Critical Remarks of Previous Works

The review of prior research demonstrates that people detection and occupancy monitoring have been approached from multiple perspectives, each offering distinct advantages and limitations. Wireless-based methods, such as Wi-Fi probe request analysis [19], [20] and Bluetooth beacon tracking [21], have proven cost-effective and scalable in large-scale deployments. Their reliance on existing infrastructure reduces hardware expenses, making them attractive for budget-constrained institutions. However, their effectiveness is inherently tied to user device-carrying behavior, meaning individuals without Wi-Fi or Bluetooth-enabled devices are excluded from counts. This dependency results in incomplete or biased occupancy data, which is particularly problematic in environments such as libraries, where not every visitor uses a connected device. Furthermore, concerns about privacy persist, as device-based tracking can unintentionally expose personal identifiers.

Hybrid and sensor fusion approaches [17], [18] provide robustness by combining multiple sensing modalities (e.g., video, Wi-Fi, and environmental data). These systems improve accuracy in challenging conditions such as low lighting or heavy occlusion and reduce the likelihood of false positives. Nevertheless, they also demand significant investment in hardware, integration expertise, and calibration. Such requirements make them less practical for medium-sized institutions like university libraries, where budgets and technical staff may be limited.

Classical computer vision methods [6], [7], [10], [11] historically formed the foundation of automated people counting. Background subtraction, optical flow, and contour analysis offered

low-cost, real-time solutions suitable for early surveillance systems. Yet, their reliance on handcrafted features makes them brittle under real-world conditions, where lighting variability, occlusion, and crowd density significantly degrade performance. These limitations explain why classical CV methods have been largely replaced by deep learning-based approaches, which can adapt more flexibly to environmental complexities.

Real-world deployments in libraries and campuses [13]–[16] further illustrate the trade-offs among sensing strategies. Thermal imaging preserves privacy but introduces high costs and reduced resolution in large open spaces. Wi-Fi sensing is scalable but inherently device-dependent, as highlighted by Wang and Shao [14]. Vision-based systems offer fine-grained spatial resolution, even at the seat level [16], but often raise privacy concerns and require substantial computational resources. Hybrid multi-sensor deployments [17], [18] achieve strong robustness, yet their infrastructure complexity limits replicability in typical academic libraries.

From these comparisons, it is evident that no single prior approach fully satisfies the requirements of accuracy, privacy, scalability, and affordability in library contexts. Most systems compromise one dimension to improve another. This gap motivates the adoption of lightweight, vision-based deep learning pipelines such as YOLOv8 with multi-object tracking. Unlike classical CV methods, modern deep learning models offer robustness against occlusion and lighting variation while maintaining real-time performance with GPU acceleration. Compared to Wi-Fi and Bluetooth-based systems, vision-based approaches count individuals directly rather than inferring presence through devices, ensuring unbiased coverage of all users. Furthermore, by combining detection with strict zone-transition logic and privacy-preserving data storage (aggregate counts without images), the proposed system balances accuracy, interpretability, and ethical concerns.

In summary, while previous works have contributed valuable insights and solutions to the problem of occupancy monitoring, their limitations underscore the need for a dedicated approach tailored to academic libraries. The proposed system builds on the strengths of prior research—accuracy, real-time capability, and historical analytics—while mitigating key weaknesses such as device dependency, infrastructure complexity, and privacy risks.

Chapter 3 System Methodology/Approach

3.1 System Design Diagram / Equation

This section presents the mathematical models and system equations used in the proposed people counting system. These equations support real-time counting, prediction, statistical analysis, and object tracking to ensure accurate monitoring and decision-making.

3.1.1 Core Counting Equations

1) Inside Count Calculation

The number of people currently inside the study corner is calculated as:

$$\begin{aligned} &\text{Inside Count Calculation:} \\ \text{Inside}(t) &= \max(0, \text{Entries}(t) - \text{Exits}(t)) \end{aligned}$$

Where:

- $\text{Entries}(t)$ = cumulative entry count at time t
- $\text{Exits}(t)$ = cumulative exit count at time t
- $\max(0, \cdot)$ ensures non-negative values

2) Peak Inside Calculation

The maximum number of people inside during a specific time period is defined as:

$$\begin{aligned} &\text{Peak Inside:} \\ \text{Peak Inside} &= \max(\text{Inside}(t)) \end{aligned}$$

Where:

- t represents the selected time period (hour, day, week, or month)

3.1.2 Prediction Algorithms

1)Next Hour Prediction

The predicted number of people inside for the next hour is computed using historical averages:

$$\text{Next Hour Prediction:} \\ \text{Predicted Inside}(h+1) = (1/n) * \Sigma \text{Inside}(h+1,i)$$

Where:

- $h+1$ = next hour
- n = number of historical samples for the same hour and day of week
- $\text{Inside}(h+1,i)$ = historical inside count of sample i

2)Prediction Confidence Score

The confidence score of the prediction is calculated as:

$$\text{Confidence} = \min(80, \max(50, 60 + (n \times 2)))$$

Where:

- n = number of historical data samples
- Confidence range: **50% – 80%**

3)Today's Remaining Entries Prediction

Remaining Entries:

$$\text{Remaining Entries} = \max(0, \text{Predicted Total} - \text{Actual So Far})$$

Where:

- Predicted Total = average end-of-day entries for the same day of week
- Actual So Far = current cumulative entries

4) Tomorrow's Total Entries Prediction

$$\text{Predicted Total(tomorrow)} = (1/m) * \Sigma \text{ Total Entries}$$

Where:

- m = number of historical days with the same day of week
- Total Entries = total entries of historical day

3.1.3 Average Calculations

1) Daily Average (Last 30 Days)

$$\text{Daily Avg} = (1/30) * \Sigma \text{ Entries}$$

$$\text{Weekly Avg} = (1/8) * \Sigma \text{ Weekly Entries}$$

$$\text{Monthly Avg} = (1/12) * \Sigma \text{ Monthly Entries}$$

3.1.4 Percentage Change Calculations

Percentage Change Between Periods

$$\text{Change \%} = ((\text{Current} - \text{Previous}) / \text{Previous}) \times 100\%$$

$$\text{Difference \%} = ((\text{Today} - \text{Average}) / \text{Average}) \times 100\%$$

3.1.5 Capacity Management

Capacity Utilization Percentage

$$\text{Capacity \%} = (\text{Inside} / \text{Max Capacity}) \times 100\%$$

Where:

- Max Capacity = 200 (configurable)
- Alert is triggered when Capacity % \geq 80%

3.1.6 Object Tracking Equations

1) Kalman Filter State Prediction

$$x_k = F x_{k-1} + w_{k-1}$$

$$P_k = F P_{k-1} F^T + Q$$

Where:

- x_k = state vector at time step k
- $x_k = [x, y, v_x, v_y]^T$ (position and velocity)
- F = state transition matrix
- P_k = error covariance matrix
- Q = process noise covariance matrix
- w_{k-1} = process noise

2) Kalman Filter Update

After receiving a new measurement, the state is updated as:

$$K_k = P_k H^T (H P_k H^T + R)^{-1}$$

$$x_k = x_k + K_k (z_k - H x_k)$$

$$P_k = (I - K_k H) P_k$$

Where:

- K_k = Kalman gain
- H = observation matrix
- R = measurement noise covariance
- z_k = measurement vector
- I = identity matrix

3)Trajectory Prediction

When an object is temporarily lost, its future position after n missing frames is estimated using average velocity:

$$\hat{x}_n = x_0 + \bar{v}_x \times n$$

$$\hat{y}_n = y_0 + \bar{v}_y \times n$$

Where:

- (x_0, y_0) = last known position
- (\bar{v}_x, \bar{v}_y) = average velocity
- n = number of missing frames

Summary

The proposed system integrates:

- **Real-time calculations** for inside count and peak detection
- **Statistical predictions** using historical averages and confidence scoring
- **Time-series aggregation** for hourly, daily, weekly, and monthly analysis
- **Object tracking models** based on Kalman filtering
- **Capacity management** through percentage-based alerts

3.2 System Architecture Diagram

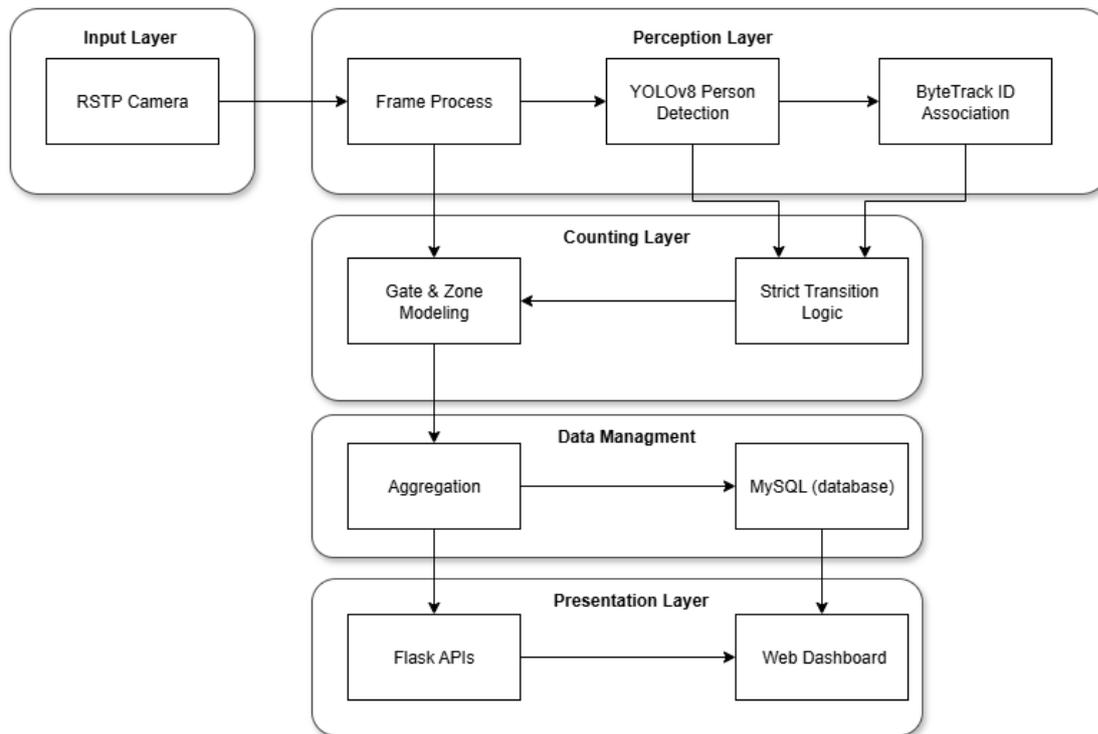


Figure 3.2 architecture diagram

Figure 3.2 illustrates the overall system architecture of the proposed people counting system. The system is designed using a layered architecture approach to ensure clear separation of responsibilities, scalability, and ease of maintenance. Each layer plays a specific role in transforming raw video input into meaningful counting information that can be visualized and analysed by users.

The system begins at the input layer, where an RTSP IP camera continuously captures and streams real-time video of the monitored area. This live video stream serves as the primary data source for the system and provides continuous visual input for subsequent processing stages.

The perception layer is responsible for visual analysis and object understanding. Video frames obtained from the camera are first processed by the frame processing module, which handles tasks such as frame buffering and pre-processing to ensure stable and consistent input. The processed frames are then passed to the YOLOv8 person detection module, where deep learning-based object detection is performed to identify human subjects in each frame. Detected persons are represented by bounding boxes with associated confidence scores. To

maintain identity consistency across frames, the ByteTrack ID association module assigns unique tracking IDs to each detected person and tracks their movement over time. This tracking mechanism helps to handle occlusion and reduces identity switching, which is essential for accurate counting.

Following detection and tracking, the system enters the counting layer, where movement information is interpreted to determine entry and exit events. Virtual gates and zones are defined within the camera view to represent different spatial regions, such as outside, gate, and inside areas. As tracked individuals move across these zones, strict transition logic is applied to validate their movement patterns. This logic ensures that only valid and complete transitions through the gate region are counted, thereby preventing double counting and reducing false detections caused by abrupt movements or tracking noise.

The data management layer is responsible for organizing and storing counting results. Counting events generated by the counting layer are aggregated into structured time-based summaries, such as per-minute, hourly, and daily statistics. These aggregated results are then stored in a MySQL database, enabling persistent storage, historical analysis, and long-term usage monitoring.

Finally, the presentation layer provides access to system outputs for end users. Flask-based application programming interfaces (APIs) retrieve real-time and historical data from the database and deliver it to the frontend. The web dashboard displays live occupancy information, historical trends, peak usage periods, and other analytical insights through an intuitive user interface. This allows users to monitor space utilization in real time and make informed decisions based on historical patterns.

Overall, the proposed system architecture enables efficient real-time people counting by integrating video acquisition, intelligent perception, reliable counting logic, robust data management, and user-friendly visualization within a unified framework.

3.3 Use Case Diagram and Description

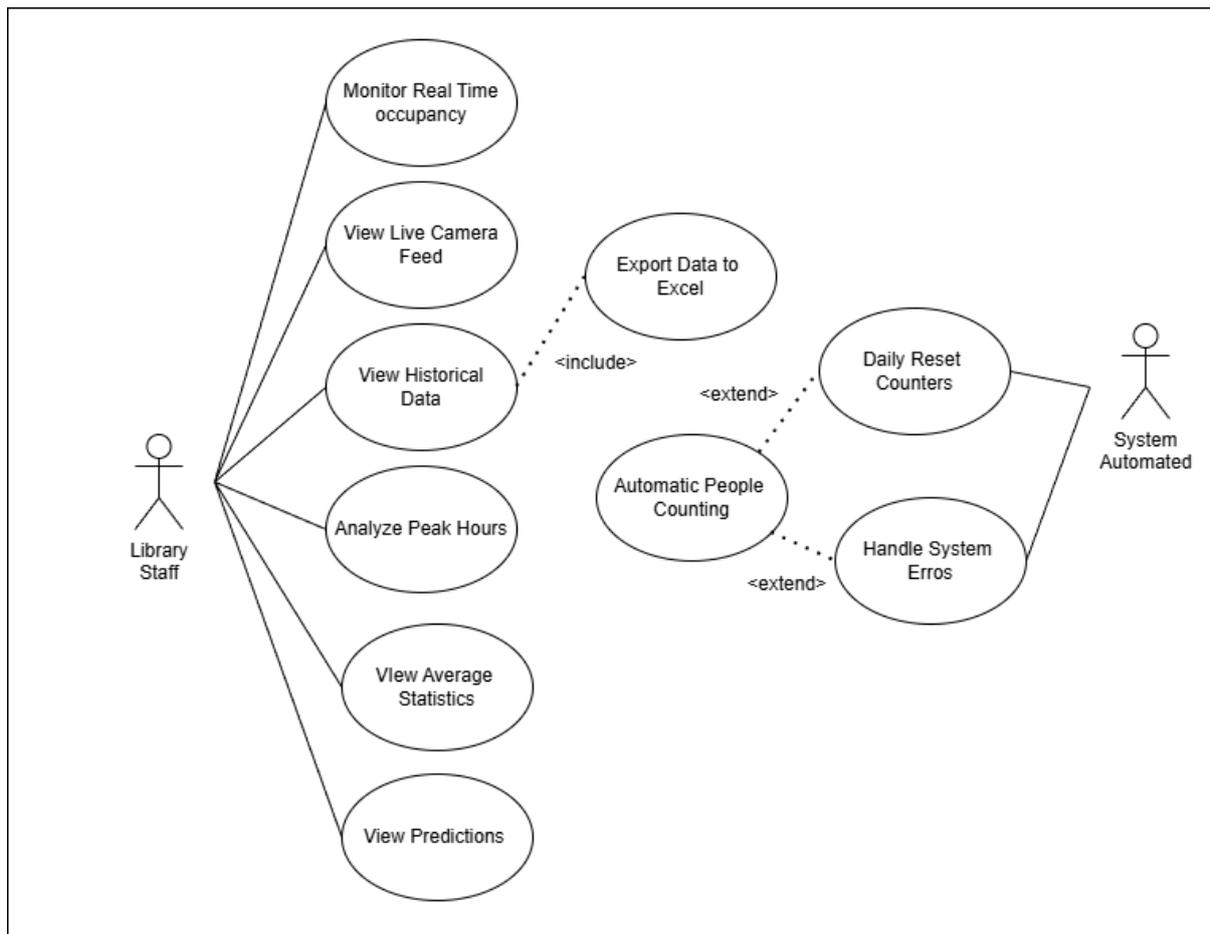


Figure 3.3: Use Case Diagram for the Proposed People Counting System

The Use Case Diagram illustrates the functional interactions between users and the proposed Library People Counting System. It provides a high-level view of how different actors interact with the system to monitor library occupancy, analyse usage patterns, and manage automated processes.

There are two main actors involved in the system: Library Staff and the System (Automated). The Library Staff represents authorized users who access the system through a web-based dashboard, while the System actor represents automated background processes that operate continuously without human intervention.

The Library Staff actor is responsible for monitoring and analysing library occupancy information. Through the Monitor Real-Time Occupancy use case, staff members can view the current number of people inside the library based on live detection results. This allows staff to ensure that occupancy levels remain within acceptable limits. In addition, staff can

access the View Live Camera Feed use case to observe the real-time video stream from the surveillance camera, which includes visual overlays such as bounding boxes generated by the detection model.

To support operational and analytical decision-making, the Library Staff can also access historical and analytical features. The View Historical Data use case allows staff to review past occupancy records stored in the database. From this historical information, staff can perform deeper analysis using the Analyse Peak Hours use case, which helps identify high-traffic periods in the library. Furthermore, the View Average Statistics use case enables staff to compare current occupancy trends against calculated daily or periodic averages, providing useful benchmarks for performance evaluation. The View Predictions use case allows staff to observe forecasted occupancy levels generated based on historical patterns, supporting proactive planning and crowd management.

The Export Data to Excel use case is included within the View Historical Data functionality. This relationship is represented using the «include» relationship, indicating that exporting data requires historical data to be retrieved first. This feature allows staff to generate structured Excel reports for administrative and record-keeping purposes.

On the automation side, the System (Automated) actor handles all background operations essential for continuous system functionality. The core automated function is Automatic People Counting, where the system continuously processes video frames, detects individuals, tracks their movements, and determines entry and exit events. This use case operates independently without user interaction.

The system also includes two extended use cases associated with automatic counting. The Daily Reset Counters use case extends Automatic People Counting by resetting daily entry and exit counters at the start of a new day while preserving historical records. Similarly, the Handle System Errors use case extends the automatic counting process by managing unexpected issues such as camera disconnections or processing errors. These extensions ensure that the system remains reliable and operates smoothly over long periods.

Overall, the Use Case Diagram demonstrates a clear separation between user-driven functionalities and automated system processes. It highlights how the Library Staff interacts with the system for monitoring and analysis purposes, while the automated components

continuously maintain accurate people counting and system stability. This design ensures efficiency, reliability, and ease of use for real-world library occupancy management.

3.4 Activity Diagram

3.4.1 People Counting Process Flow

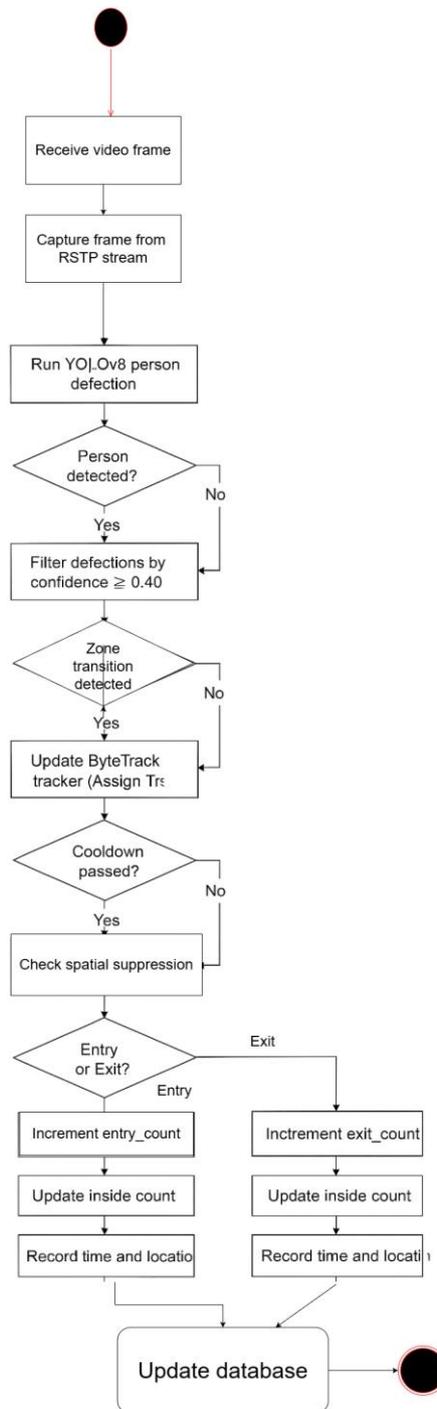


Figure 3.4.1 Activity Diagram for People Counting Process Flow

The People Counting Process Flow activity diagram illustrates the step-by-step operational workflow of the proposed system in detecting, tracking, and counting people entering and exiting the study corner in real time. This process begins when the system continuously receives video frames from the RTSP camera stream.

Initially, each video frame is captured and resized to a fixed resolution of 1280×720 pixels to ensure consistent input quality and to optimize processing performance. The resized frame is then passed to the YOLOv8 object detection model, which identifies human objects within the frame. Only detections classified as persons with a confidence score of 0.40 or higher are considered valid. In addition, detections with bounding box heights smaller than 50 pixels are discarded to reduce false positives caused by distant or partially visible objects.

Once valid detections are obtained, the system updates the ByteTrack tracker to maintain consistent identities for detected individuals across consecutive frames. ByteTrack assigns or updates unique tracking IDs and applies Kalman filtering to predict object movement, allowing the system to handle temporary occlusions, or missed detections. Appearance histogram matching is also applied to improve re-identification accuracy.

For each tracked individual, the system determines the current spatial zone, which can be categorized as outside, gate, or inside. The system then analyses the recent zone history over the last three frames to detect valid zone transitions. Several validation checks are applied to prevent double counting, including a cooldown period of 0.5 seconds, spatial suppression within a 30-pixel radius over 12 frames, and a minimum track age requirement of five frames.

If a valid entry transition is detected, the system increments the entry counter and updates the current inside count. Similarly, if a valid exit transition is identified, the exit counter is incremented, and the inside count is decremented accordingly. For each counting event, the system records the corresponding time and location for logging and analysis purposes.

After processing counting events, the system updates the database when a count event occurs. This ensures that real-time data is preserved while minimizing excessive database writes. The system then prepares the processed frame for visualization by drawing bounding boxes, zone boundaries, and updated count information before storing the frame for API access and dashboard display.

Finally, the activity cycle completes by proceeding to the next video frame, allowing the system to operate continuously in real time. Through this structured workflow, the activity diagram demonstrates how the system ensures accurate, reliable, and efficient people counting under real-world conditions.

3.4.2 Dashboard Data Retrieval Activity Diagram

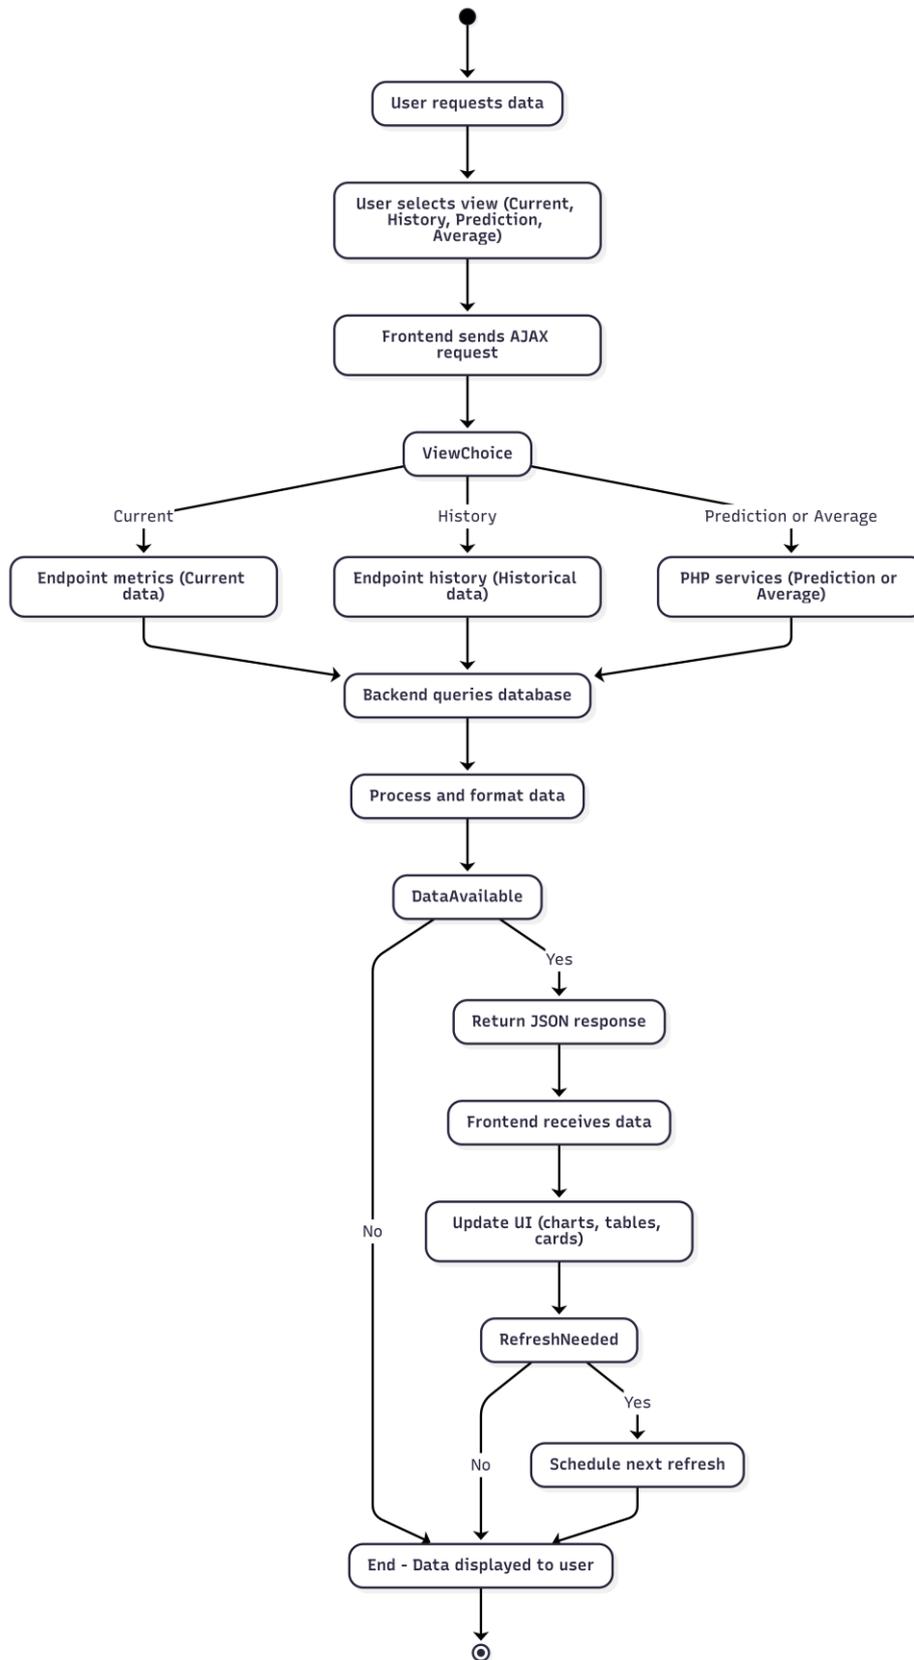


Figure 3.4.2 People Counting Process Flow Activity Diagram

This activity diagram represents the process of retrieving and displaying dashboard data based on the user's selected view. The process starts when the user requests data from the dashboard.

Initially, the user selects a view, which can be Current, History, Prediction, or Average. After the selection is made, the frontend sends an AJAX request to the backend system. The request is then evaluated to determine which view has been selected.

Based on the selected view, the request is routed to the appropriate backend service. Requests for current data are handled by the metrics endpoint, historical data requests are sent to the history endpoint, and prediction or average data requests are processed by dedicated PHP services. All of these services query the database to retrieve the required information.

Once the data is retrieved, the backend processes and formats it into a suitable structure. A decision point is then reached to check whether the requested data is available. If the data is not available, the system returns an error or an empty response, and the process ends.

If the data is available, the backend returns the processed data in JSON format to the frontend. The frontend receives the response and updates the user interface by refreshing charts, tables, and cards to display the new data.

Finally, the system determines whether a refresh is needed, such as in real-time monitoring scenarios. If a refresh is required, the system schedules the next refresh cycle. Otherwise, the process ends with the data successfully displayed to the user.

3.4.3 Daily Reset Process Flow

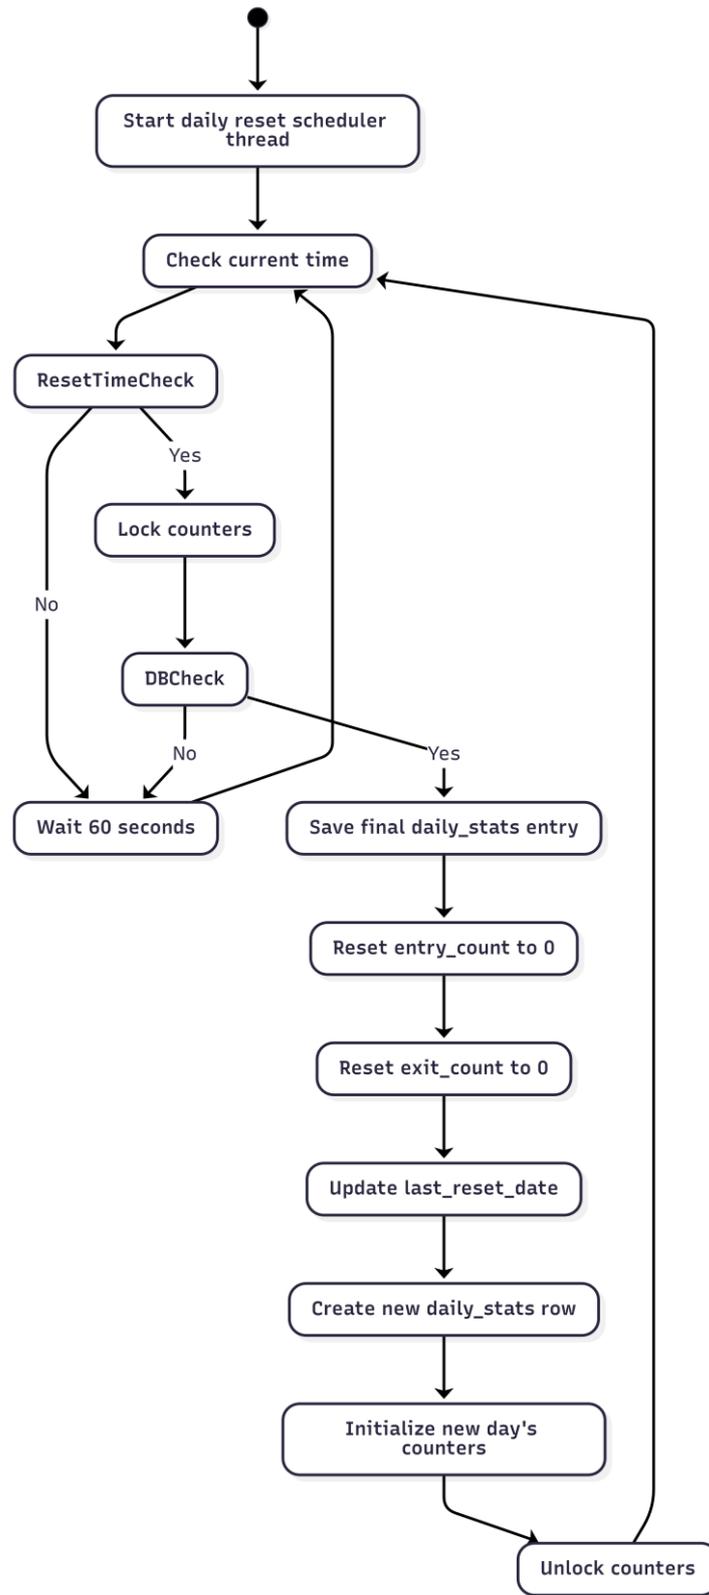


Figure 3.4.3 Daily Reset Process Flow Activity Diagram

This activity diagram illustrates the workflow of the system's daily reset scheduler, which runs continuously to manage daily statistics.

The process begins with the scheduler thread, which periodically checks the current system time. A decision is then made to determine whether it is reset time (midnight).

If it is not yet midnight, the scheduler waits for 60 seconds before checking the time again. This loop continues until the reset time is reached.

When it is midnight, the scheduler first locks the counters to prevent concurrent modifications during the reset process. It then checks whether the database connection is available. If the connection is not available, the scheduler waits and retries in the next cycle. If the database is available, the system proceeds with the reset operations.

The reset operations include:

1. Saving the final `daily_stats` entry for the previous day.
2. Resetting `entry_count` and `exit_count` to zero.
3. Updating the `last_reset_date` to reflect the new day.
4. Creating a new `daily_stats` row for the new day.
5. Initializing counters for the new day.

After completing these steps, the scheduler unlocks the counters to allow normal operations to resume. The process then loops back to continuously check the time, ensuring that the daily reset occurs every day at midnight.

Overall, this diagram emphasizes the automated, continuous nature of the scheduler, the use of decision points for time and database availability, and the **critical locking mechanism** to maintain data consistency during the reset.

3.4.4 Predictive Analytics Calculation Flow

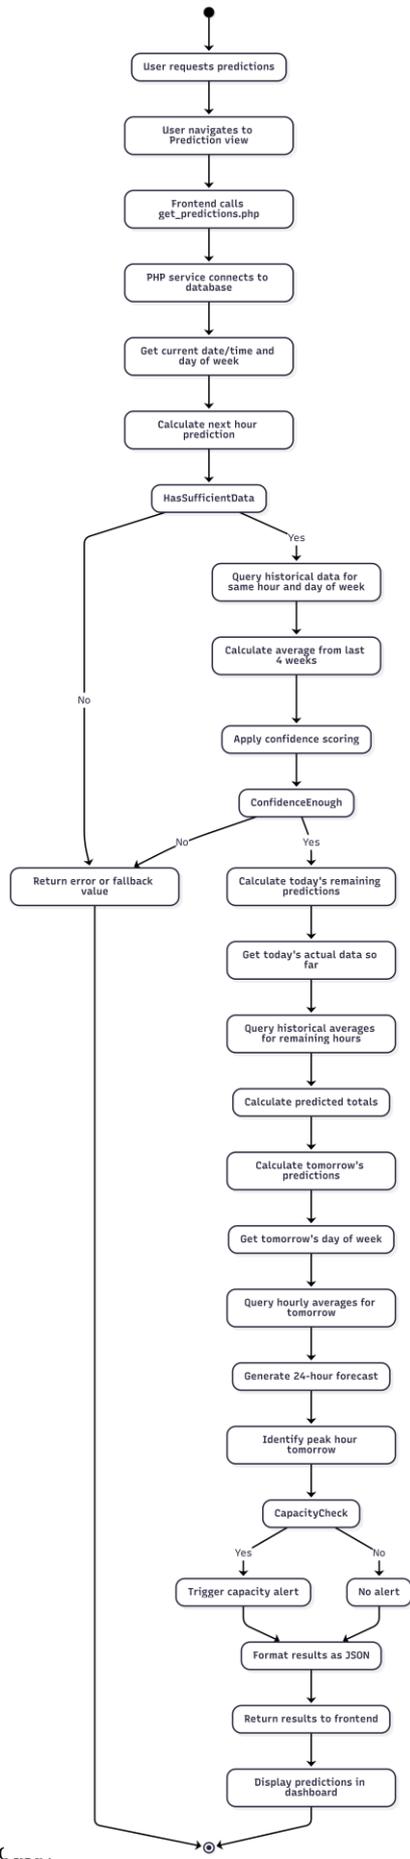


Figure 3.4.4 Predictive Analytics Calculation Flow activity diagram

This activity diagram represents the workflow for calculating and displaying predictive analytics on the dashboard. The process starts when the user requests predictions by navigating to the Prediction view. The frontend then calls the `get_predictions.php` PHP service to retrieve the data.

The PHP service first connects to the database and obtains the current date, time, and day of the week. It then calculates the next hour prediction by querying historical data for the same hour and day of the week over the past four weeks, averaging the results, and applying a confidence scoring mechanism. A decision point checks whether there is sufficient historical data. If not, a fallback or error is returned. Another decision verifies whether the calculated confidence meets the threshold; if it does not, the system also returns a fallback response.

If the confidence threshold is met, the service proceeds to calculate today's remaining predictions. It retrieves the actual data collected so far for today, queries historical averages for the remaining hours, and calculates predicted totals. Next, it calculates tomorrow's predictions by determining the day of the week, querying hourly averages, and generating a 24-hour forecast. The system also identifies the peak hour for tomorrow.

Another decision checks whether the predicted usage exceeds the system's capacity threshold (e.g., 80%). If it does, a capacity alert is triggered. Finally, all results are formatted as JSON and returned to the frontend, where they are displayed on the dashboard. The process ends once the predictions are successfully displayed.

This diagram highlights the step-by-step predictive calculations, the decision points for data availability, confidence scoring, and capacity alerts, and the flow of information between the frontend, PHP service, and database.

Chapter 4: System Design

4.1 System Block Diagram

This section describes the overall system structure and data flow of the proposed People Counter System. Top-down design diagrams are used to illustrate how data is captured, processed, stored, and presented to users. The diagrams provide a clear understanding of system functionality before detailed component specifications are discussed in later sections

4.1.1 Overall System Architecture Block Diagram

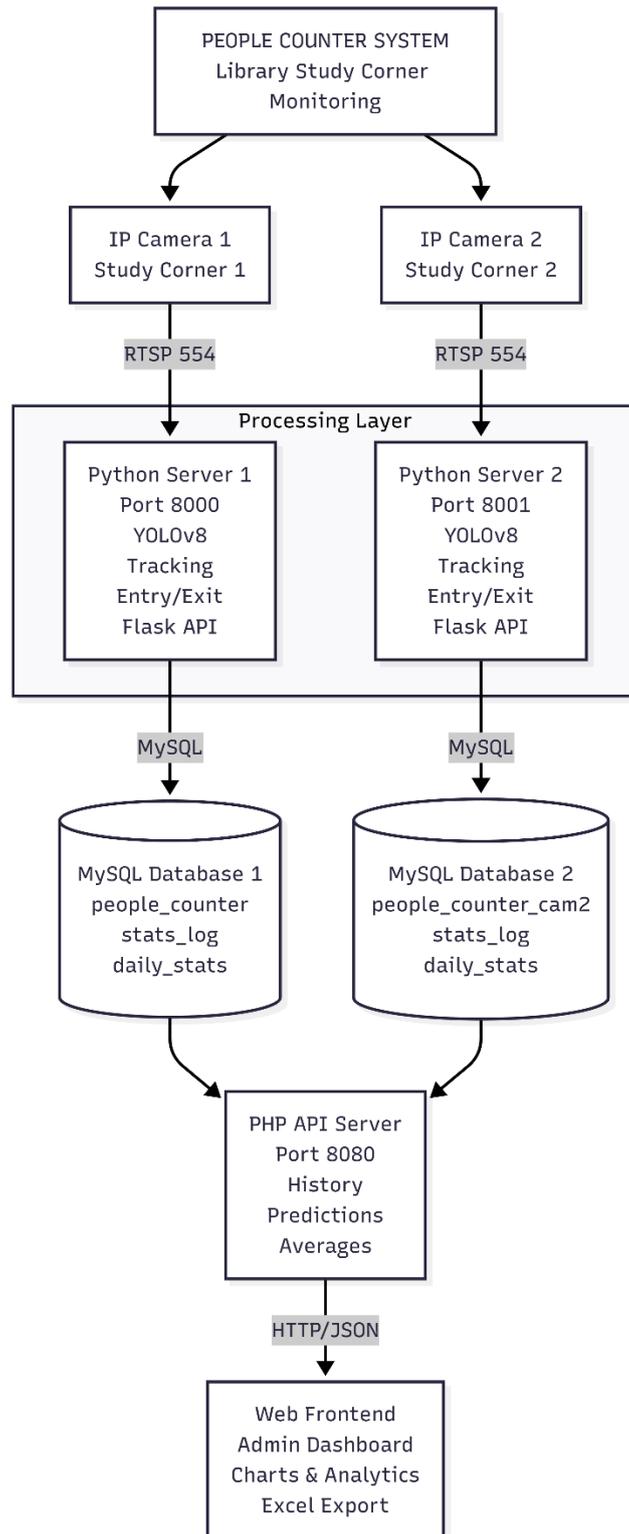


Figure 4.1.1 System Block Diagram

Figure X illustrates the overall system architecture of the proposed People Counter System for Library Study Corners. The system is designed using a multi-layer and multi-camera architecture to enable real-time people counting, data storage, and visualization.

At the input layer, two IP cameras are installed at different study corners within the library. Each camera continuously captures video footage of the entrance area and streams the video to the processing server using the Real-Time Streaming Protocol (RTSP). This allows real-time transmission of video data over the local area network.

The processing layer consists of two independent Python-based processing servers, where each server is dedicated to a single camera. Each Python server receives the RTSP video stream and performs real-time video frame processing. The YOLOv8 object detection model is used to detect human objects in each frame, while an object tracking algorithm assigns a unique tracking ID to each detected person. Based on predefined virtual zones, the system applies entry and exit detection logic to determine whether a person is entering or leaving the study corner. The processing servers also expose a Flask-based API to provide real-time metrics and video streams to other system components.

The data layer is implemented using MySQL databases, where each camera has its own dedicated database. This design ensures independent operation and avoids data conflicts between cameras. The databases store both timestamped logs and aggregated daily statistics, enabling efficient retrieval of historical and real-time occupancy data.

To support data access and analytics, a PHP-based API server is used as an intermediate layer between the databases and the web interface. The PHP API executes structured SQL queries to retrieve historical records, compute statistical averages, and generate prediction-related information. All responses are returned in JSON format, which ensures compatibility with web-based applications.

Finally, the presentation layer consists of a web-based frontend accessed through a standard web browser. The frontend displays real-time camera streams, live occupancy statistics, historical data visualizations, and analytical charts. It also provides functionality for exporting data in spreadsheet format for reporting and analysis purposes.

Overall, this architecture adopts a modular and scalable design, where each camera-processing pipeline operates independently while sharing a unified frontend interface. This approach improves system reliability, simplifies maintenance, and allows additional cameras to be integrated into the system with minimal changes.

4.1.2 Data Flow Diagram

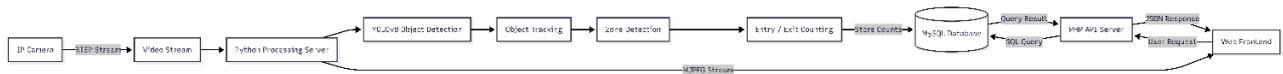


Figure 4.1.2 Data Flow Diagram

Figure 4.1.2 illustrates the data flow within the proposed People Counter System. The system operates through three main flows: real-time processing, data retrieval, and video streaming, each of which is critical for accurate monitoring and reporting of library occupancy.

Real-time Processing Flow: Video data is captured by the IP cameras and transmitted to the Python processing server as a live stream. Each video frame is analyzed by the YOLOv8 object detection model to detect human objects. Detected objects are then tracked using the object tracking module, which maintains persistent IDs for individuals across frames. Using predefined zone boundaries, the system determines entry and exit events. The counted results are periodically stored in the MySQL database for long-term persistence and analytics.

Data Retrieval Flow: When a user requests information from the web interface, the request is sent to the PHP API server. The server executes structured SQL queries on the database and returns the retrieved information in JSON format. The frontend then displays the data in tables, charts, and statistical summaries, allowing administrators to monitor library usage efficiently.

Video Streaming Flow: For real-time monitoring, the Python processing server retrieves video frames from the cameras, encodes them as MJPEG streams, and exposes them through Flask endpoints. The web frontend accesses these endpoints to display live camera feeds directly in the browser, enabling users to observe occupancy in real-time.

This data flow diagram provides a high-level overview of how data moves and is processed, demonstrating the system's capability to capture, analyze, store, and present people-counting information effectively.

4.1.3 Software Architecture Layers

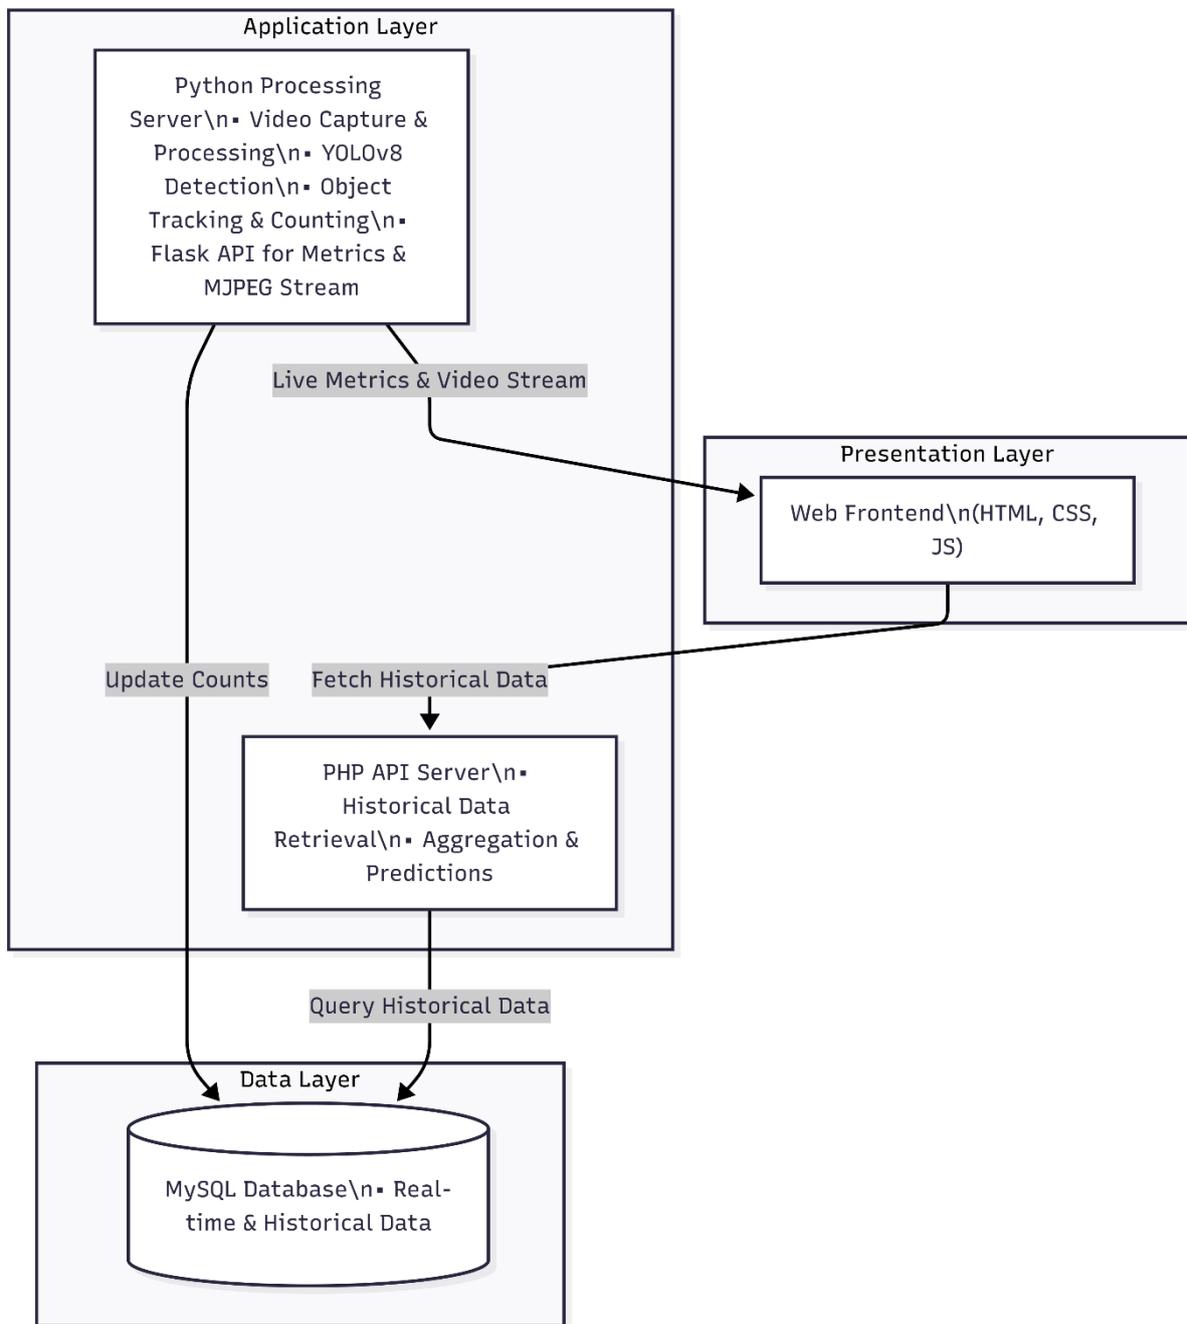


Figure 4.1.3 three-tier software architecture diagram

Software Architecture Diagram Explanation

The proposed system follows a **three-tier architecture**: Presentation Layer, Application Layer, and Data Layer. The updated diagram illustrates how data flows through the system from real-time detection to frontend visualization.

1. Application Layer – Python Processing Server

The Python server is the core of the system and acts as the first step in processing:

- It captures video frames from IP cameras in real time via RTSP streams.
- The YOLOv8 object detection model identifies human objects in each frame.
- Object tracking (using ByteTrack and Kalman filter) assigns persistent IDs and follows individuals across frames.
- Entry and exit logic determines when a person enters or leaves the predefined zones.
- Counts of entries, exits, and people inside are stored in memory and periodically written to the MySQL database.
- The Python server also exposes a Flask API to provide live metrics and a MJPEG video stream directly to the frontend.

2. Data Layer – MySQL Database

The MySQL database stores both real-time metrics and historical data:

- Python server writes live counts and updates daily statistics.
- PHP API server queries the database for historical data and analytics.
- This separation ensures data persistence, reliability, and efficient querying for frontend displays and reports.

3. Application Layer – PHP API Server

The PHP API server handles:

- Historical data retrieval from the database.
- Aggregations and predictions for analytics dashboards.
- It serves JSON responses to the frontend for rendering charts, tables, and statistics.

4. Presentation Layer – Web Frontend

The **web frontend** provides the user interface:

- Displays real-time metrics and live video streams received directly from the Python server.

- Fetches historical data and predictions from the PHP API server.
- Presents information in dashboards, charts, and tables for monitoring and analysis.

4.2 System Components Specifications

4.2.1 Hardware Components

This section describes the physical devices used in the People Counter System. The main hardware includes IP cameras for capturing real-time video of the study corners and a processing server to analyze the video streams. Each component is selected to ensure reliable data acquisition, accurate people counting, and seamless integration with the system's software modules.

4.2.1.1 IP Cameras

Component	Specification / Detail
Type	IP Network Camera with RTSP support
Resolution	Minimum 1280×720 (HD)
Protocol	RTSP (Real-Time Streaming Protocol)
Network	Ethernet connection
Power Supply	PoE (Power over Ethernet) or AC adapter
Location	Camera 1: Study Corner 1 Camera 2: Study Corner 2
IP Address	Camera 1: 192.168.246.65 Camera 2: 192.168.246.66
RTSP Port	554
Authentication	Username/Password (admin/admin123)
Functions	- Capture real-time video feed- Stream video to processing server- Support people

	counting and tracking- Enable accurate monitoring and occupancy logging
--	---

Table 4.2: IP Camera Specifications and Functions

Table 4.2 summarizes the key specifications and functions of the IP cameras used in the People Counter System for the library study corners. These cameras capture high-definition video of the entrances, stream the footage in real time to the processing server via RTSP protocol, and integrate with the people counting algorithm to monitor occupancy accurately. The detailed hardware specifications, such as resolution, network type, power options, and IP configuration, ensure reliable performance and scalability for future expansions..

4.2.2 Software Components

4.2.2.1 Python Processing Server

Component	Specification / Detail
Language	Python 3.8+
Framework	Flask (Web framework)
AI Model	YOLOv8x (Ultralytics)
Tracking Algorithm	ByteTrack
Libraries	<ul style="list-style-type: none"> - OpenCV (cv2) - Video processing - PyTorch - Deep learning - NumPy - Numerical operations - PyMySQL - Database connectivity - Ultralytics YOLO - Object detection

Table 4.2.2.1.1: Python Processing Server Specifications

Server	Camera	Port	Database	RTSP URL
Server 1	Camera 1	8000	people_counter	rtsp://admin:admin123@192.168.246.65:554/. ..
Server 2	Camera 2	8001	people_counter_cam 2	rtsp://admin:admin123@192.168.246.66:554/. ..

Table 4.2.2.1.2: Server Configuration

Key Features and Functionality:

The Python Processing Server is responsible for real-time people detection and tracking using YOLOv8x and the ByteTrack algorithm. Each server is assigned to a specific camera and processes its RTSP video stream independently. Key features include:

- **Real-time object detection:** Detects only the person class in each video frame.
- **Multi-object tracking:** Maintains consistent IDs across frames using Kalman filtering to handle occlusions and trajectories.
- **Zone-based entry/exit detection:** Detects when a person moves through defined zones (outside → gate → inside) to accurately count entries and exits.
- **Automatic daily reset:** Counters reset at midnight, ensuring daily statistics are separated.
- **Database persistence:** Updates MySQL database every 60 seconds to log stats.
- **MJPEG video streaming:** Provides live video feed to the frontend for monitoring.
- **RESTful API endpoints:** Enables easy data access and integration with frontend or external applications.

Endpoint	Method	Description
/metrics	GET	Returns current counts (entries, exits, inside) in JSON format
/stream	GET	Provides MJPEG video stream for live display
/frame.jpg	GET	Returns a single JPEG frame for snapshot purposes

Endpoint	Method	Description
/history	GET	Optional endpoint to retrieve historical data

Table 4.2.2.1.3: Python Server API Endpoints

Explanation:

The Python Processing Server serves as the core processing layer of the People Counter System. Each server runs independently for its assigned camera, capturing video via RTSP, performing detection and tracking, and updating the MySQL database. The Flask framework provides lightweight API services for frontend visualization, while the YOLOv8x and ByteTrack combination ensures accurate and reliable people counting. This modular design allows easy scalability and maintenance, supporting multiple cameras without interference.

4.2.2.2 PHP API Server

Component	Specification / Detail
Programming Language	PHP 7.4+
Web Server	Apache (via XAMPP)
Port	8080
Database Driver	PDO (MySQL)
Data Format	JSON
Role	Acts as middleware between MySQL database and web frontend

Table 4.2.2.2.1: PHP API Server Specifications

API File	Description
get_history_data.php	Retrieves historical people-counting data based on selected period (hourly, daily,

	weekly, monthly), supports pagination, and returns results in JSON format
get_predictions.php	Generates predictions using historical patterns, including next-hour occupancy, remaining entries for today, next-day forecast, and capacity alerts
get_averages.php	Computes statistical averages such as daily, weekly, and monthly averages, day-of-week patterns, hourly averages, and comparisons between current values and historical averages
get_history_data_cam2.php	Provides camera 2–specific historical data with the same functionality as get_history_data.php

Table 4.2.2.2.2: PHP API Files and Functions

Key Features and Functionality

The PHP API Server functions as the application-layer middleware of the People Counter System. It handles client requests from the web frontend, performs database queries and statistical calculations, and returns structured data in JSON format. Key functionalities include historical data retrieval, prediction generation based on time-based patterns, and statistical analysis across different periods. The use of PDO ensures secure and efficient database communication, while Apache and XAMPP provide a stable deployment environment.

The PHP API Server separates data access and business logic from the frontend, improving maintainability and scalability. By centralizing database queries and analytics within the PHP layer, the system ensures consistent data processing for multiple cameras and views. This design also allows future extensions, such as additional analytics modules or authentication mechanisms, without modifying the frontend or Python processing servers.

4.2.2.3 MySQL Database

Component	Specification / Detail
Database System	MySQL 5.7+ / MariaDB 10.3+
Port	3306
Character Set	utf8mb4
Storage Engine	InnoDB
Primary Role	Persistent storage for people counting statistics

Table 4.2.2.3.1: MySQL Database Specifications

Database Structure

Two separate databases are implemented to support independent operation of each camera while maintaining identical schemas for consistency and scalability.

Column Name	Data Type	Description
id	INT (PK, AUTO_INCREMENT)	Unique record identifier
ts	DATETIME	Timestamp of the record
day	DATE (Generated)	Date extracted from timestamp
entries	INT	Number of entries
exits	INT	Number of exits
inside	INT	Current number of people inside

Table 4.2.2.3.2: stats_log Table Structure

Column Name	Data Type	Description
day	DATE (PK)	Date of record
entries	INT	Total entries for the day
exits	INT	Total exits for the day
inside	INT	Current occupancy

peak_inside	INT	Maximum occupancy recorded
open_time	DATETIME	Time first entry occurred
close_time	DATETIME	Time last exit occurred

Table 4.2.2.3.3: daily_stats Table Structure

Database 2: people_counter_cam2 (Camera 2)

The second database follows the **same table structure and schema** as people_counter but operates independently to store data for Camera 2. This separation prevents data conflicts and allows each camera stream to be processed and analyzed independently.

Database Operations

The MySQL databases support both real-time and historical data storage with the following operations:

- **INSERT:** Executed every 60 seconds to log data into the stats_log table.
- **UPDATE:** Performed in real time to maintain aggregated values in the daily_stats table.
- **SELECT:** Triggered on demand by PHP API queries for data visualization, analysis, and reporting.
- **Automatic Daily Reset:** At 00:00 each day, counters are reset and a new daily record is initialized.

The MySQL database serves as the data layer of the People Counter System, providing reliable and structured storage for both real-time and historical occupancy data. The use of the InnoDB storage engine ensures transaction safety and data integrity, while indexed columns optimize query performance for analytics and reporting. Maintaining separate databases for each camera enhances modularity, simplifies maintenance, and supports future system scalability.

4.2.2.4 Web Frontend

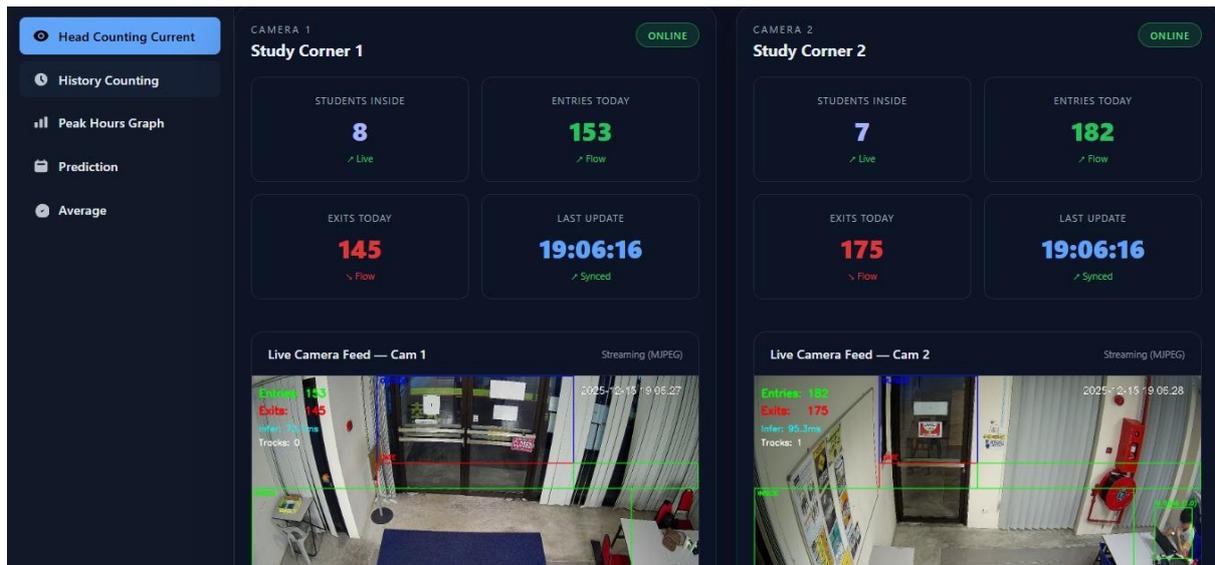


Figure 4.2.2.4.1 Dashboard

Component	Specification / Detail
Programming Languages	HTML5, CSS3, JavaScript (ES6+)
Visualization Library	Chart.js
Export Library	XLSX.js
Browser Support	Google Chrome, Mozilla Firefox, Microsoft Edge
Primary Role	Data visualization and system monitoring interface

Table 4.2.2.4.2: Web Frontend Specifications

Page Name	Description
admin_dashboard_clean.html	Main administrative dashboard providing real-time camera feeds, live statistics, navigation sidebar, peak hour analysis, prediction results, and average analysis

history_counting.html	Historical data page for Camera 1, featuring tabular views, interactive charts, and Excel export functionality
history_counting_cam2.html	Historical data page for Camera 2 with the same features as the Camera 1 history page

Table 4.2.2.4.3: Web Frontend Pages

Key Features

The Web Frontend provides an intuitive and interactive interface for monitoring and analyzing people-counting data. Its key features include:

- Real-time data updates with a 1-second refresh interval
- Interactive charts for visualizing trends and patterns
- Date and period selection for customized historical analysis
- Pagination support for handling large datasets efficiently
- Excel export with well-formatted tables for reporting purposes
- Responsive design for consistent usability across different screen sizes

The Web Frontend serves as the presentation layer of the People Counter System, enabling administrators to monitor real-time occupancy, review historical trends, and analyze predictions through a web browser. By integrating Chart.js for data visualization and XLSX.js for data export, the interface supports both operational monitoring and decision-making tasks. The separation between real-time dashboards and historical analysis pages improves usability, while the responsive design ensures accessibility on various devices and display resolutions.

4.2.3.1 YOLOv8 Model

Component	Specification / Detail
Model Variant	YOLOv8x (Extra Large)
Input Size	640 × 640 pixels
Detection Class	Person only (Class 0)
Confidence Threshold	0.40
Framework	PyTorch (Ultralytics)
Execution Device	GPU (CUDA) / CPU

Table 4.2.3.1.1: YOLOv8 Model Specifications

Metric	Value
Inference Speed	Approximately 30–50 ms per frame (GPU)
Detection Accuracy	High precision for human detection
Model Size	Approximately 130 MB

Table 4.2.3.1.2 Performance Characteristics

YOLOv8x is used as the primary object detection model due to its high accuracy and robustness in detecting human subjects in real-time video streams. By restricting detection to the person class only, the system reduces false positives and improves inference efficiency. GPU acceleration using CUDA enables real-time processing suitable for continuous monitoring in library study corners.

4.2.3.2 Object Tracking

Component	Specification / Detail
Tracking Algorithm	ByteTrack
Tracker Configuration	botsort.yaml
Tracking Type	Multi-object tracking
ID Management	Persistent object IDs across frames
Occlusion Handling	Supported
Trajectory Prediction	Enabled

Table 4.2.3.2.1: Object Tracking Specifications

ByteTrack is employed to associate detected persons across consecutive video frames by maintaining consistent tracking IDs. It effectively handles occlusion scenarios and missed detections, ensuring accurate tracking continuity. The use of the botsort.yaml configuration improves association robustness, which is essential for reliable entry and exit counting.

4.2.3.3 Kalman Filter

Component	Specification / Detail
Filter Type	4-state Kalman Filter
State Vector	[x, y, vx, vy]
Update Rate	Per frame (~30 FPS)
Primary Purpose	Position prediction and smoothing

Table 4.2.3.3.1: Kalman Filter Specifications

The Kalman Filter is integrated into the tracking pipeline to predict object positions between frames and smooth trajectory estimates. By modeling both position and velocity, the filter compensates for temporary detection losses and reduces noise in object movement. This

enhances tracking stability and contributes to accurate people counting, especially in crowded or partially occluded environments.

Together, YOLOv8x, ByteTrack, and the Kalman Filter form a robust detection and tracking pipeline. YOLOv8x provides accurate person detection, ByteTrack ensures consistent multi-object tracking, and the Kalman Filter enhances motion prediction and stability, enabling reliable real-time people counting in the library study corners.

4.2.3.4 YOLOv8 Training and Fine-Tuning Process

Objective of Model Fine-Tuning

The objective of the YOLOv8 fine-tuning process is to adapt a pre-trained object detection model to the specific environment of the library study corners. Although the original YOLOv8 model is trained on large-scale general datasets, variations in camera angle, lighting conditions, background layout, and crowd density can reduce detection accuracy in real-world deployment. By fine-tuning the model using data captured directly from the study corners, the system achieves higher person-detection accuracy and improved people-counting reliability.

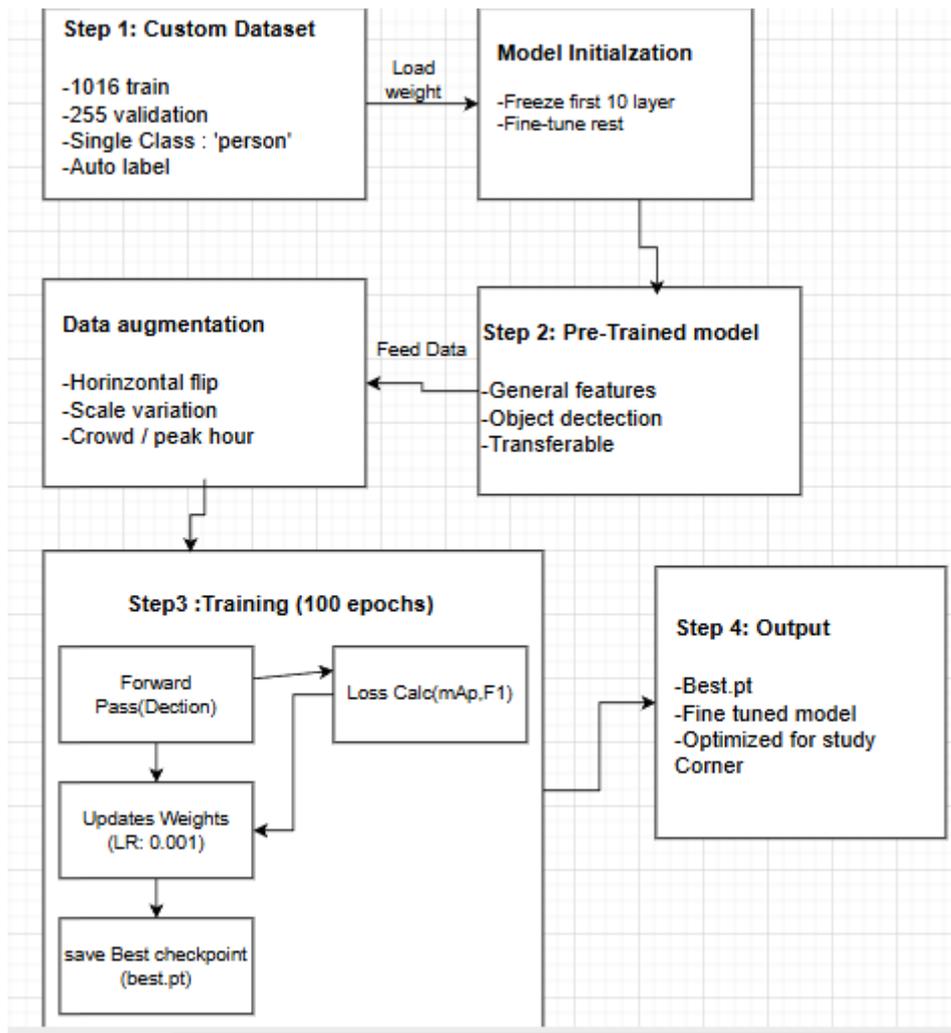


Figure 4.2.3.4.1 Flow of the fine tuning process

Figure X illustrates the workflow of the YOLOv8 fine-tuning process applied in the People Counter System to adapt a pre-trained model to the library study-corner environment.

The process begins with Step 1: Custom Dataset Preparation, where a domain-specific dataset is created using images extracted from the RTSP camera streams. A total of 1,016 images are used for training and 255 images for validation. Since the system focuses solely on people counting, only a single class (“person”) is defined. Images are automatically labeled and further refined to ensure accurate bounding boxes.

In Step 2: Model Initialization, a pre-trained YOLOv8 model is loaded with existing weights learned from large-scale datasets. To preserve general visual features such as edges and shapes, the first 10 backbone layers are frozen, while the remaining layers are fine-tuned.

This transfer learning strategy allows the model to adapt efficiently to the study-corner environment with limited training data.

Before training, data augmentation is applied to improve generalization. Augmentation techniques include horizontal flipping, scale variation, and sampling images from different crowd densities and peak-hour scenarios. These augmentations help the model handle variations in lighting, viewpoint, and occupancy levels commonly found in real-world library settings.

In Step 3: Training, the model is trained for 100 epochs. During each epoch, a forward pass is performed to detect persons in the input images. The detection results are evaluated using loss functions based on mAP and F1 score, and model weights are updated using a learning rate of 0.001. Throughout training, the system continuously monitors performance on the validation set and saves the best-performing checkpoint (best.pt).

Finally, Step 4: Output produces a fine-tuned YOLOv8 model (best.pt) that is optimized for the study-corner environment. This model demonstrates improved detection accuracy, reduced false positives, and more stable performance compared to the original pre-trained model. The optimized model is then deployed in the live Python processing servers for real-time people counting.

Dataset Preparation

Frame Extraction from RTSP Streams

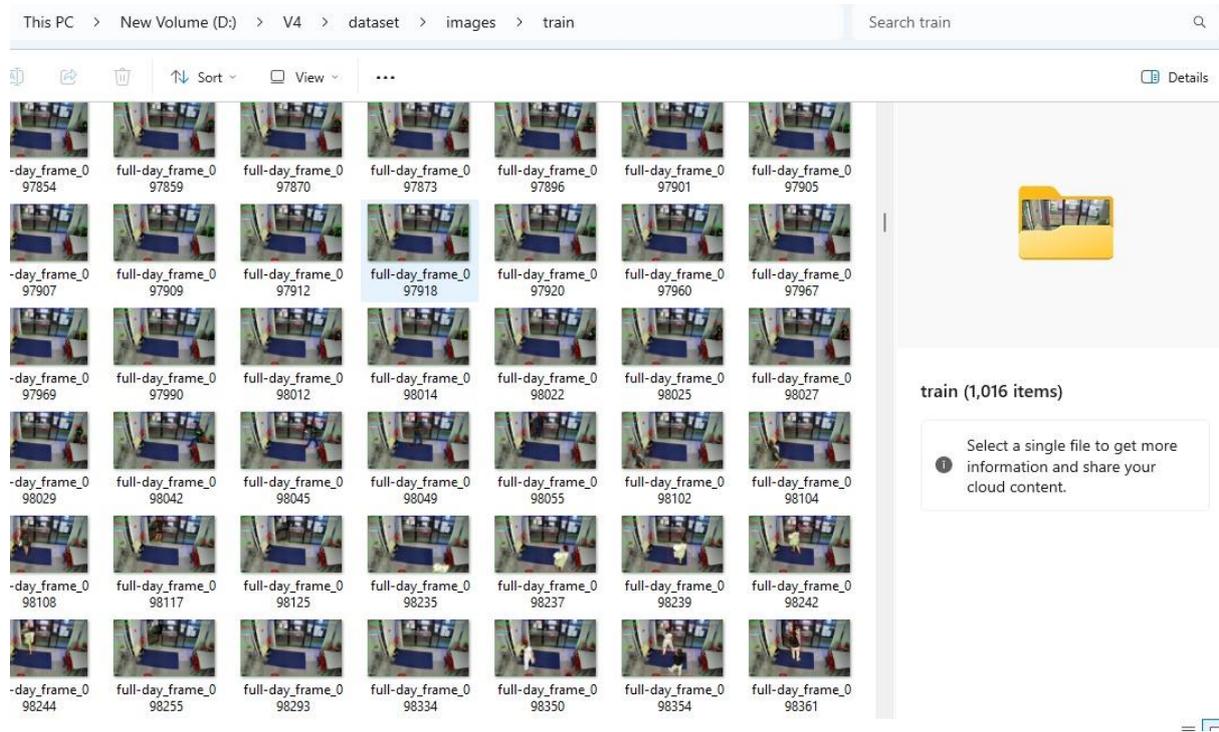


Figure 4.2.3.4.2 Data set extract from recording

Training data was obtained by extracting image frames from RTSP video streams of the installed IP cameras. Video recordings were captured from both Study Corner 1 and Study Corner 2 at different times of the day to ensure diverse lighting conditions and crowd densities. A custom script (`extract_training_frames.py`) was used to extract frames at fixed intervals from the recorded videos. This approach ensured that the training dataset closely reflected the actual operational environment of the People Counter System.

Image Annotation in YOLO Format

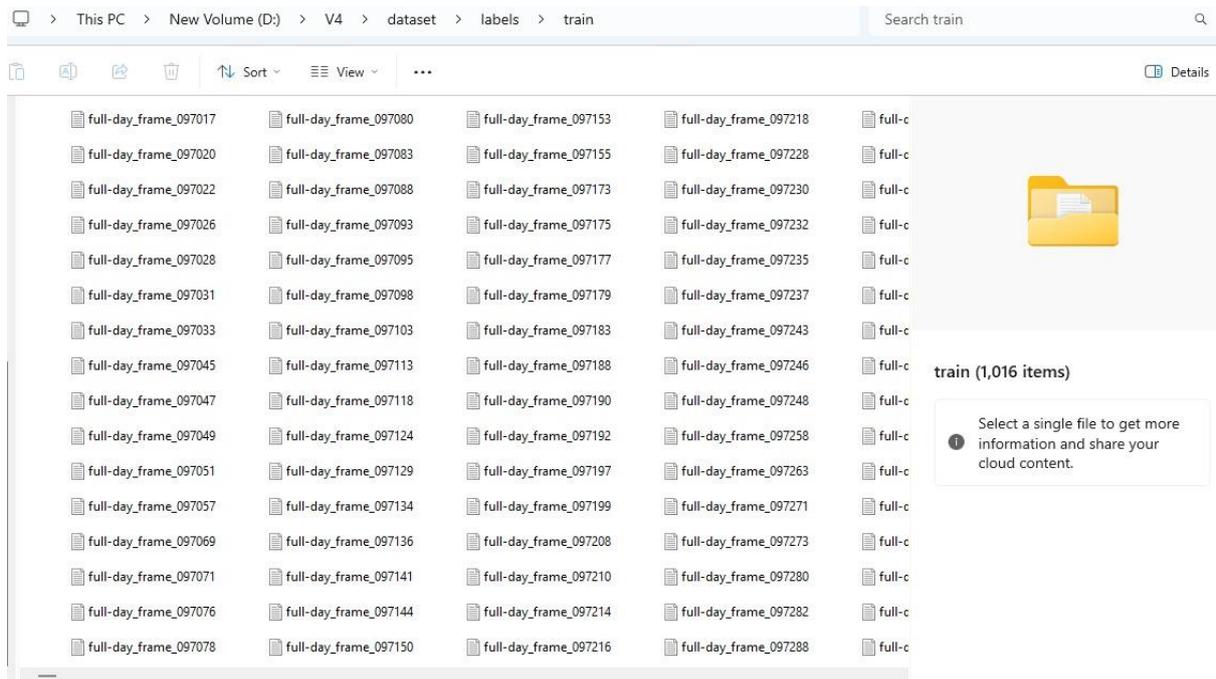


Figure 4.2.3.4.3 Labeled data

All extracted frames were first auto-annotated using the `auto_label_frames.py` script, which applies a pre-trained YOLO model to detect people and generate YOLO-format label files automatically. For each image, the script creates a corresponding text file containing normalized bounding box coordinates, where class 0 represents a person (the only class used, since the system focuses exclusively on people counting). After auto-labeling, the generated annotations were manually reviewed and corrected (adding missed detections, removing false positives, and adjusting bounding boxes where necessary) to ensure high-quality training data.

Dataset Organization and Configuration

The annotated dataset was divided into training and validation subsets to enable performance evaluation during training. Images and labels were organized into structured directories (`images/train`, `images/val`, `labels/train`, and `labels/val`). A `dataset.yaml` file was created to define dataset paths and class information. Care was taken to balance the dataset by including scenes with different occupancy levels, camera perspectives, and lighting conditions from both study corners.

Transfer Learning and Training Pipeline

Base Model Selection

The YOLOv8x (extra-large) model was selected as the base model due to its superior detection accuracy. Transfer learning was applied by initializing training with pre-trained weights (yolov8x.pt), allowing the model to reuse general visual features learned from large datasets while adapting higher-level features to the study-corner environment. For systems with limited GPU memory, smaller variants such as YOLOv8l or YOLOv8m can be used.

Training Execution

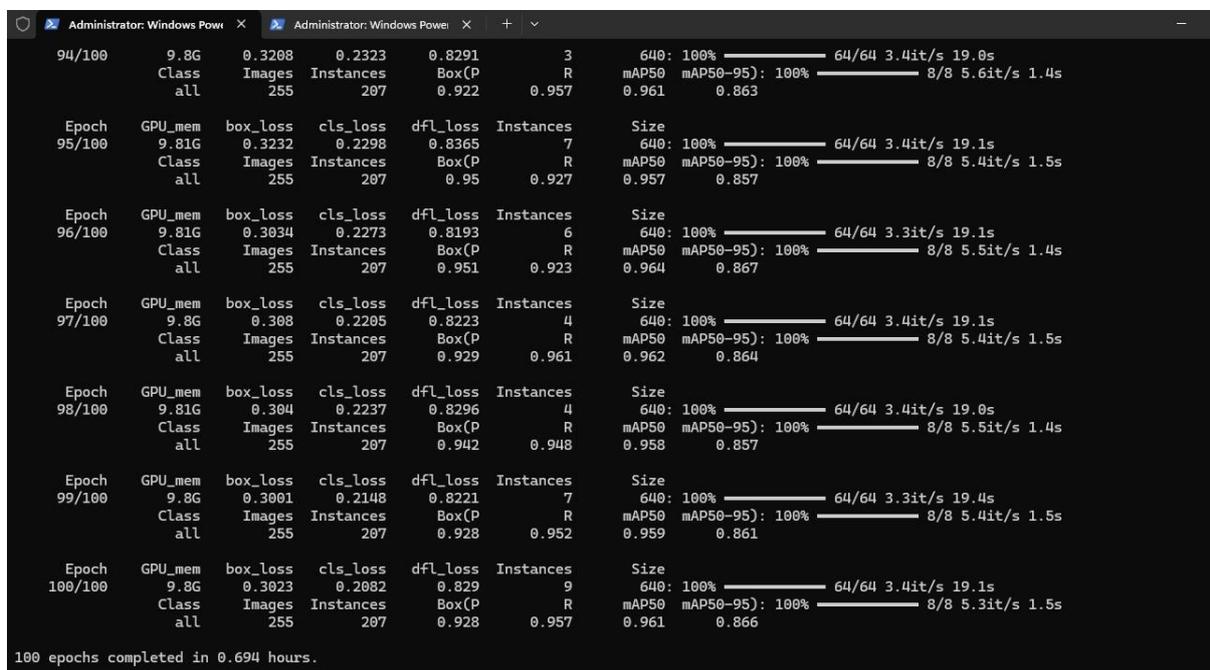


Figure 4.2.3.4.4 Training process

Model training was conducted using a custom training script (train_yolo.py) that wraps the Ultralytics YOLO training API. The training process was executed for up to 100 epochs with an input image size of 640×640 pixels. A batch size of 16 and an initial learning rate of 0.001 were selected to balance convergence stability and GPU resource usage. The first 10 backbone layers were frozen during initial training to prevent overfitting and preserve general feature representations.

Training performance was monitored using validation metrics such as precision, recall, mAP@50, and mAP@50–95. Early stopping with a patience value of approximately 50 epochs was applied to prevent over-training.

Hyperparameter Configuration

Key hyperparameters were selected based on YOLOv8 fine-tuning best practices and hardware constraints. The confidence threshold during inference was set to 0.40 to reduce false detections while maintaining sensitivity. A moderate batch size and conservative learning rate were used to ensure stable training. These parameters were adjusted as necessary to fit available GPU memory and achieve smooth convergence.

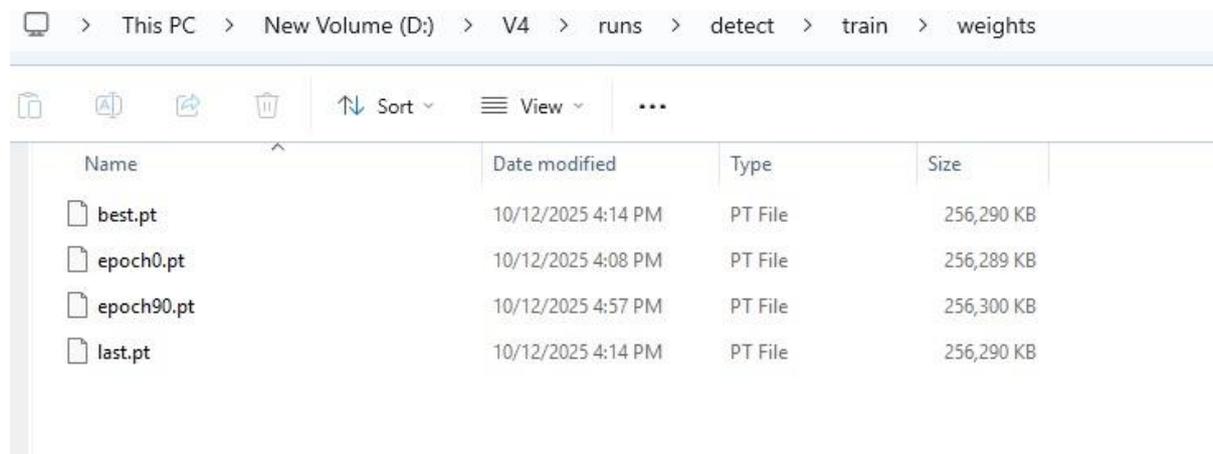
Data Augmentation Strategy

Light data augmentation techniques were applied to improve model generalization. These included horizontal flipping, image scaling, HSV color jittering, and small-angle rotations. Aggressive augmentations such as heavy blur were avoided, as they negatively impact the detection of small or partially occluded persons in surveillance footage.

Hardware Requirements for Training

Training was performed on a system equipped with an NVIDIA CUDA-enabled GPU with at least 8GB of VRAM, which is recommended for YOLOv8x. The system also utilized a modern multi-core CPU, 8–16GB of RAM, and SSD storage to support efficient data loading and checkpoint saving. In cases of limited GPU memory, batch size reduction or smaller YOLOv8 variants may be used.

Training Outputs and Deployment



Name	Date modified	Type	Size
best.pt	10/12/2025 4:14 PM	PT File	256,290 KB
epoch0.pt	10/12/2025 4:08 PM	PT File	256,289 KB
epoch90.pt	10/12/2025 4:57 PM	PT File	256,300 KB
last.pt	10/12/2025 4:14 PM	PT File	256,290 KB

Figure 4.2.3.4.5 Best.pt

Upon completion of training, the best-performing model weights were saved as best.pt, while the final training checkpoint was saved as last.pt. Additional outputs, including loss curves and mAP plots, were automatically generated by the Ultralytics framework. The fine-tuned best.pt model was deployed by updating the Python processing servers to load the new weights, replacing the generic pre-trained model.

Live system verification was conducted by running inference on RTSP streams and visually comparing detection results before and after fine-tuning.

Evaluation and Acceptance Criteria

Model evaluation was conducted at both offline and system levels. Offline evaluation used validation metrics such as mAP, precision, and recall verifying improvements over the base model. Online evaluation was performed by deploying the fine-tuned model in the live People Counter System and comparing system-generated counts with manual counts during test periods. The fine-tuned model was considered acceptable when it demonstrated improved detection stability, reduced false counts, and achieved a counting accuracy of approximately 95% or higher under real-world conditions.

Through transfer learning and fine-tuning, the YOLOv8 model was successfully adapted to the library study-corner environment. By training on camera-specific data, applying appropriate augmentations, and deploying the optimized model in the live system, the People

Counter System achieved improved detection accuracy and more reliable occupancy counting compared to the baseline pre-trained mode

4.2.3.5 Average Analysis

The **Average Analysis** module provides statistical insights on library usage over various time scales. Its primary objective is to support effective planning, staffing, and capacity management by understanding patterns of entries, exits, and occupancy levels. Data for these analyses are obtained from two main sources: the `daily_stats` table, which contains aggregated daily entries, exits, and peak inside counts, and the `stats_log` table, which records raw hourly logs used to calculate derived averages.

The system computes several metrics using `get_averages.php`:

- **Daily Average (last 30 days):**
 - Average daily entries, exits, and occupancy.
 - Maximum daily entries, exits, and peak inside.
 - Percentage change compared to the previous week's average.
- **Weekly Average (last 8 weeks):**
 - Total entries and exits per week, averaged across weeks.
 - Percentage change compared to the last month's weekly averages.
- **Monthly Average (last 12 months):**
 - Monthly totals and averages of entries, exits, and peak occupancy.
- **Day-of-Week Average:**
 - Average entries, exits, and inside counts for each day of the week (Monday–Sunday).
 - Used to populate the “Average by Day of Week” bar chart.
- **Hourly Average (current day-of-week):**
 - Average entries and occupancy per hour (0–23) for the same weekday.
 - Used for the “Hourly Average” line chart.
- **Today vs Average:**

- Compares today's entries, exits, and peak occupancy against current daily averages.
- Highlights percentage differences to indicate above- or below-average usage.

Frontend Usage:

- Populates Average cards (daily, weekly, monthly, peak).
- Feeds charts such as Average by Day of Week, Hourly Average, and Monthly Average Trends.
- Provides contextual insights, e.g., "+12% vs last week," to highlight trends and support decision-making.

4.2.3.6 Prediction Feature

The **Prediction Feature** aims to forecast short-term and next-day library usage, enabling users to avoid peak hours and assisting library management in capacity planning. This functionality is implemented through `get_predictions.php`.

Inputs:

- Camera selection (cam1 or cam2), mapped to the respective databases (people_counter or people_counter_cam2).
- Current date, time, and day-of-week (derived on the server).

Historical Window:

- Typically uses the last 4 weeks of stats_log data for the same day-of-week.

Predicted Metrics:

- **Next Hour Prediction:**
 - Predicted number of people inside for the next hour.
 - Computed as the historical average at the same hour, with priority on the same day-of-week.

- Includes predicted inside count, reference entries/exits, and a confidence score.
- **Today's Remaining Hours:**
 - Predicted incremental entries and exits per hour.
 - Predicted occupancy for each remaining hour.
 - Outputs include remaining_entries, remaining_exits, predicted_total_entries, and predicted_total_exits.
 - Comparison against historical daily totals.
- **Tomorrow's Forecast:**
 - Uses all historical days with the same day-of-week as tomorrow.
 - Outputs include predicted total entries, predicted total exits, and predicted peak occupancy.
 - Provides an hourly breakdown of entries, exits, and inside counts.
- **Peak Hour Prediction:**
 - Identifies the hour tomorrow with the highest predicted occupancy.
 - Returns the peak time and expected number of people.
- **Capacity Alerts:**
 - Configurable maximum capacity (e.g., 200 people).
 - Calculates capacity percentage for each predicted hour and flags hours where occupancy $\geq 80\%$.
 - Returns a list of alert hours with time, percentage, and predicted people count.

Frontend Usage:

- Populates Prediction cards such as:
 - "Next Hour Prediction"
 - "Today's Remaining Entries"
 - "Today's Predicted Total"
 - "Today's Actual So Far"
 - "Peak Time Tomorrow"
 - "Tomorrow's Total Entries"
 - "Tomorrow's Peak Inside"
 - "Capacity Alert"
- Drives line charts:

- Today's Remaining Hourly Predictions
- Tomorrow's Hourly Predictions

Benefits:

- Helps users select less crowded times to visit the study corner.
- Supports proactive resource and capacity planning for library staff.
- Provides interpretable forecasts rather than just raw numbers.

4.3 Circuits and Components Design

4.3.1 Network Architecture

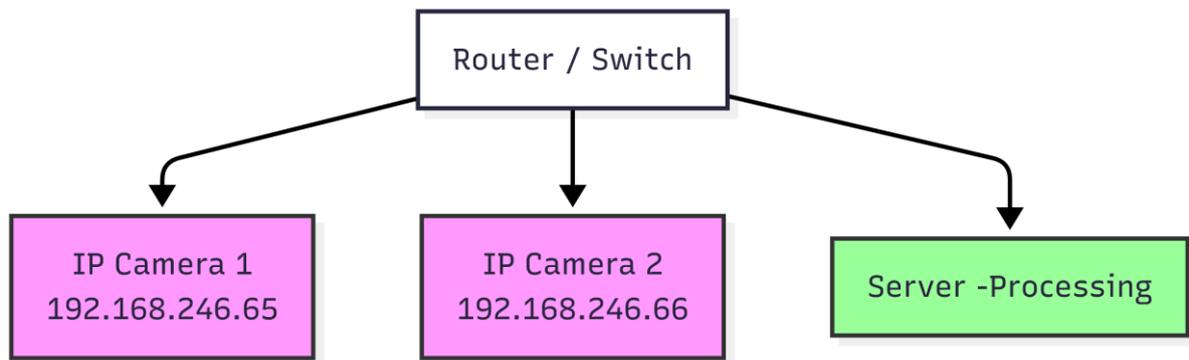


Figure 4.3.1.1: Network Architecture of the People Counter System

The network architecture of the People Counter System is designed to connect multiple IP cameras to a central processing server through a router or switch, enabling real-time video streaming and data processing. The cameras are assigned static IP addresses within the subnet 192.168.246.0/24 to ensure reliable communication, while the server handles the processing of captured frames for people counting and analytics.

Key configuration details include:

- **Subnet:** 192.168.246.0/24, allowing devices within the same network to communicate efficiently.
- **Gateway:** The router IP, which connects the local network to external networks if needed.

- **DHCP:** Enabled for automatic IP assignment, simplifying device setup.
- **Ports Used:**
 - **554:** RTSP protocol for live camera streaming.
 - **8000 & 8001:** Flask servers for backend processing.
 - **8080:** Apache/PHP server for frontend and API services.
 - **3306:** MySQL database server for storing historical logs and analytics.

The diagram visually represents the topology, showing the router/switch at the center, connected to two IP cameras and the processing server. This layout ensures that all devices can transmit video streams and analytics data seamlessly for real-time monitoring and reporting.

4.3.2 Data Flow Architecture

Real-time Processing Circuit:

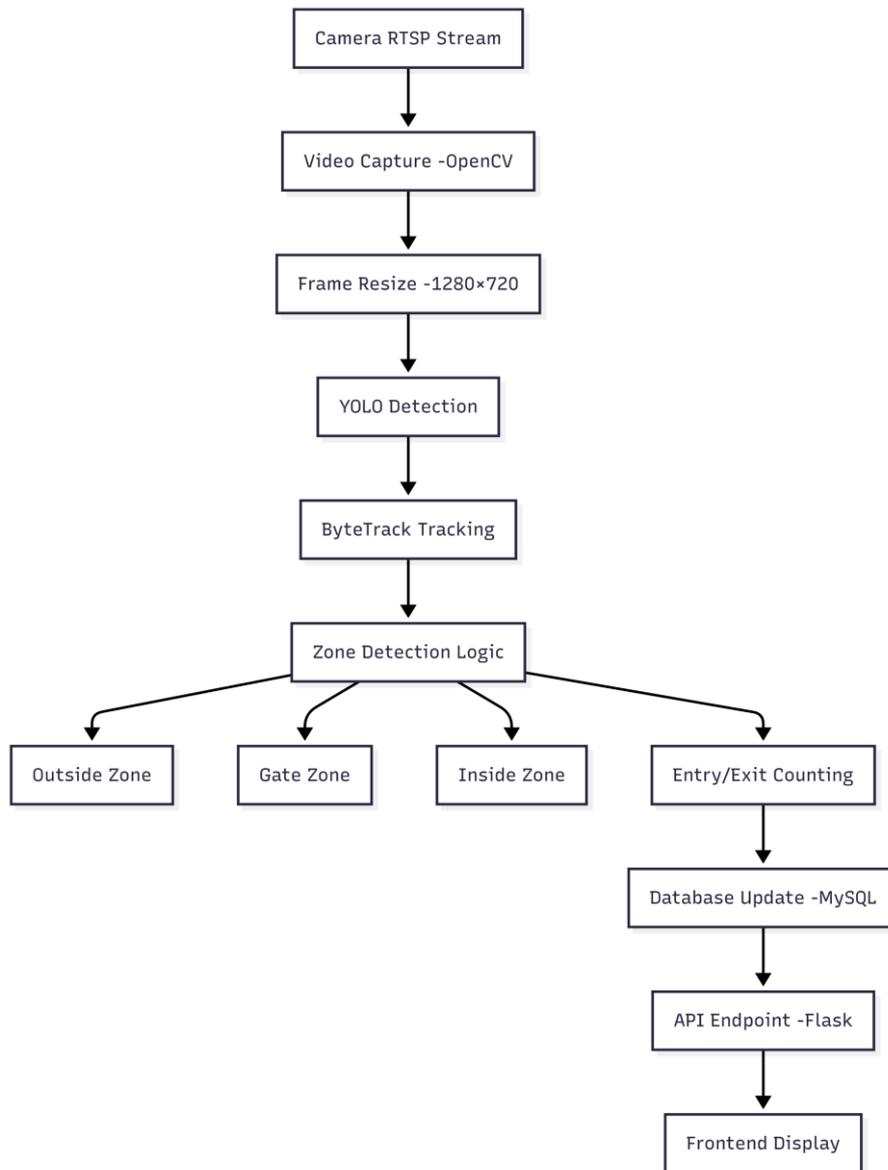


Figure 4.3.2.1 Data Flow Architecture of the Real-Time People Counter System

The data flow architecture illustrates the step-by-step process for real-time people counting in the library study corner. The system captures live video streams from IP cameras via the RTSP protocol and processes each frame using OpenCV. Each frame is resized to 1280×720 for optimized performance before being analyzed by the YOLO object detection model to identify people.

Detected objects are then tracked using ByteTrack, which assigns consistent IDs across frames to maintain tracking even during occlusions. The tracked data is passed to the zone detection logic, which categorizes people into Outside, Gate, or Inside zones. This information feeds into the entry/exit counting module, which updates the MySQL database in real time.

Finally, the processed data is exposed through a Flask API, which enables the frontend display to visualize current occupancy, entries, and exits. This architecture ensures continuous, accurate monitoring of library usage and supports further statistical analysis and predictions.

4.4 System Components Interaction Operations

4.4.1 Real-Time Processing Interaction

The real-time processing interaction describes how the system captures video, detects people, tracks their movement, and updates the database to maintain accurate occupancy counts. Two Python scripts are used to manage the camera streams:

- **Camera 1:** v6.py (Flask on port 8000, database people_counter)
- **Camera 2:** v7_cam2.py (Flask on port 8001, database people_counter_cam2)

Video Capture and Processing

The system performs the following operations:

1. Camera Initialization:

- Each Python script connects to the camera's RTSP stream.
- A background thread (ThreadedVideoReader) is started to continuously read frames.
- Video stream buffers are initialized to store incoming frames.

2. Frame Processing Loop:

Each frame goes through a continuous processing loop:

- Read frame from buffer
- Resize frame to 1280×720
- Run **YOLO detection** to identify people

- Assign IDs using **ByteTrack tracking**
- Determine the person's zone (**Outside, Gate, or Inside**)
- Evaluate **entry/exit logic** based on zone sequence
- Update counters accordingly
- Save a **database snapshot** every 60 seconds
- Display or stream the processed frame to the frontend

3. Entry/Exit Detection Logic:

- Entry is counted when a person passes through zones in order: **Outside** → **Gate** → **Inside**
- Exit is counted when moving in reverse: **Inside** → **Gate** → **Outside**
- **Cooldown period:** 0.5 seconds to avoid double-counting
- **Minimum box height:** 50 pixels to filter false detections

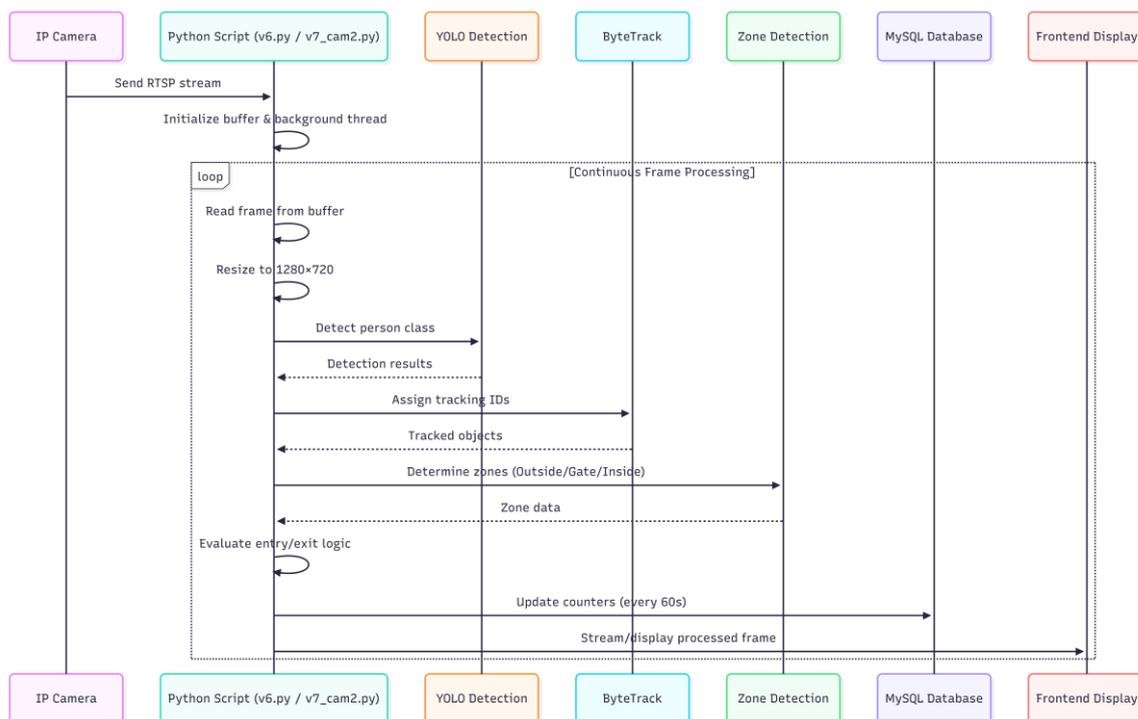


Figure 4.4.1.1: Sequence Diagram of Real-Time Processing Interaction

The sequence diagram illustrates the step-by-step interaction between system components in the real-time people counting process. The IP camera streams video frames to the Python

script (v6.py or v7_cam2.py), which initializes a buffer and background thread to continuously capture frames without delays.

Each frame undergoes preprocessing, including resizing to 1280×720 pixels for optimized performance. The preprocessed frame is then passed to the YOLO detection module, which identifies people in the frame. Detected objects are tracked across successive frames using ByteTrack, ensuring consistent IDs even when temporary occlusions occur.

The zone detection logic categorizes each tracked person into one of three zones: Outside, Gate, or Inside. This zone information is evaluated according to the entry/exit rules: a person moving from Outside → Gate → Inside is counted as an entry, and Inside → Gate → Outside is counted as an exit. To avoid false counts, a cooldown period of 0.5 seconds and a minimum box height of 50 pixels are enforced.

Processed results are periodically stored in the MySQL database (every 60 seconds) for analytics and historical records. Simultaneously, the processed frames and occupancy information are streamed to the frontend display, providing users and library staff with real-time visualization of entries, exits, and current occupancy.

This sequence ensures that video capture, detection, tracking, zone classification, and database updates operate seamlessly in a continuous loop, enabling accurate and reliable real-time monitoring of library usage.

4.4.2 Database Interaction

The database interaction module manages both read and write operations to ensure accurate recording and retrieval of occupancy data for the library study corner.

Write Operations:

Every 60 seconds, the system performs a write operation to update the database. This involves inserting the latest frame's occupancy data into the stats_log table and updating the daily_stats table to reflect cumulative entries, exits, and peak occupancy. Each update is executed as a single transaction to maintain consistency, while the system is designed to continue processing even if an error occurs during a write operation, ensuring that real-time monitoring is not interrupted.

Read Operations:

- On **system startup**, the current counters for the day are loaded from the `daily_stats` table to resume accurate tracking.
- When the system is **reset**, it queries the database for the current date and resets the counters accordingly.
- **API queries** are executed on-demand, using `SELECT` statements to retrieve historical or real-time data for analytics, visualizations, or frontend display.

This approach ensures that all occupancy and movement data are reliably recorded, available for real-time monitoring, and ready for statistical analysis and predictive features, while maintaining robustness against transient errors during database writes.

4.4.3 Frontend-Backend Interaction

4.4.3.1 Real-time Metrics Update

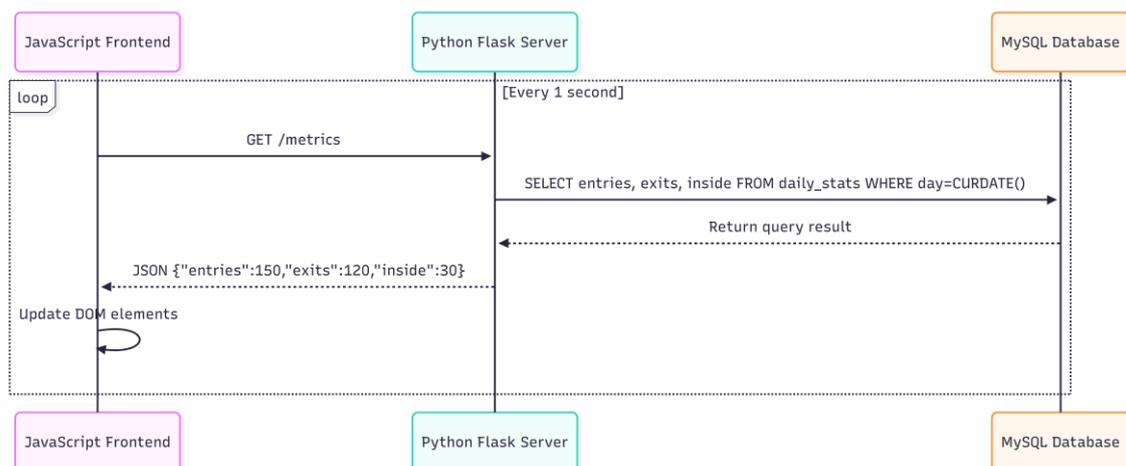


Figure 4.4.3.1.1: Real-Time Metrics Update Flow Between Frontend and Backend

The real-time metrics update illustrates how the frontend interface continuously receives and displays occupancy data from the backend. The frontend, implemented in JavaScript, periodically sends a request to the Flask server every second (`setInterval(1000ms)`) to fetch the latest metrics.

Upon receiving the request, the Flask server queries the **MySQL database**, selecting the current day's entries, exits, and inside counts from the `daily_stats` table. The server then

returns this data as a **JSON response**, for example: {"entries": 150, "exits": 120, "inside": 30}.

The frontend parses the JSON response and updates the corresponding **DOM elements** dynamically, ensuring that users see up-to-date information about library occupancy, entries, and exits in real time.

This interaction provides a seamless connection between the backend data processing and the frontend visualization, enabling responsive and accurate monitoring of the study corner.

4.4.3.2 Video Streaming Interaction

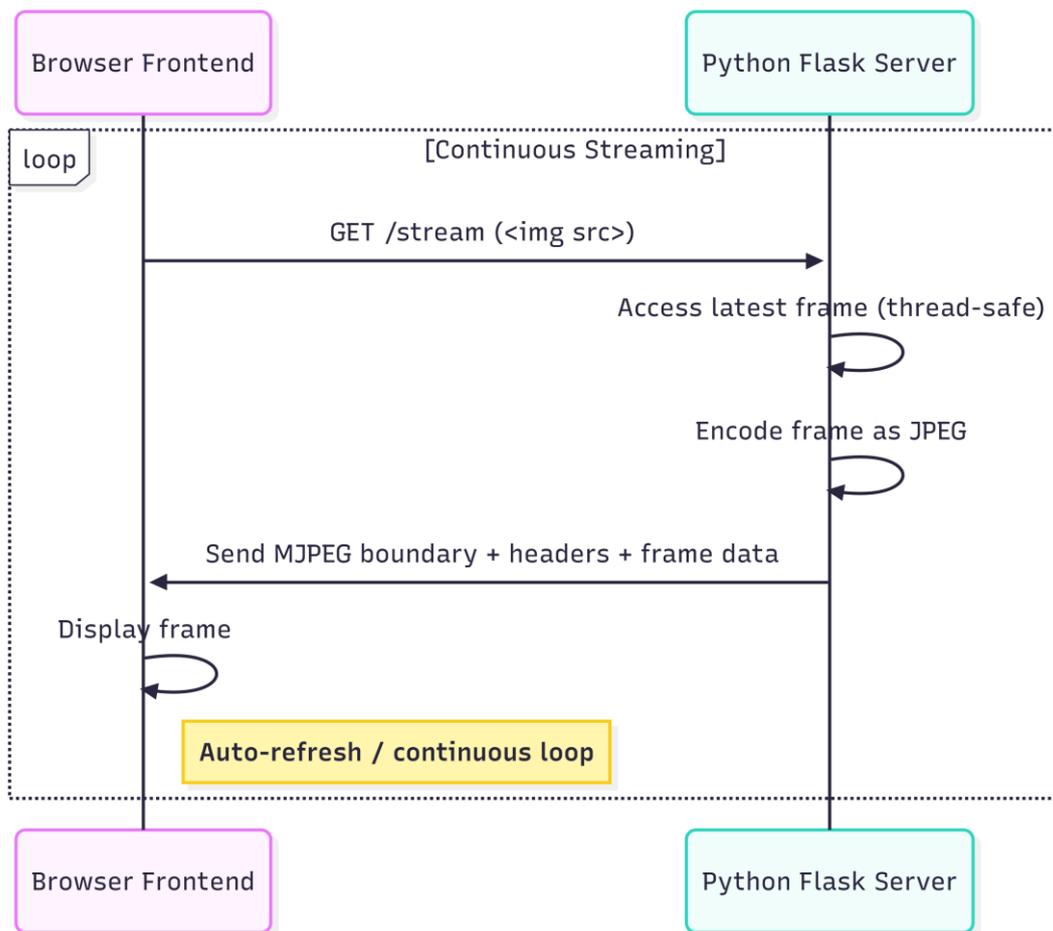


Figure 4.4.3.2.1: Video Streaming Interaction Between Frontend and Backend

This sequence diagram illustrates the interaction between the frontend and backend for real-time video streaming. The frontend embeds an `` tag pointing to the Flask server endpoint `/stream`. The Flask server continuously retrieves the latest frame from the video capture buffer using a thread-safe lock to avoid concurrency issues.

Each frame is then JPEG encoded and sent to the frontend as part of an MJPEG stream, including boundaries and headers to separate consecutive frames. The browser automatically displays each incoming frame, creating a smooth live video effect. This loop repeats continuously, ensuring that the frontend always shows the most recent frame from the camera.

4.4.3.3 Historical Data Retrieval

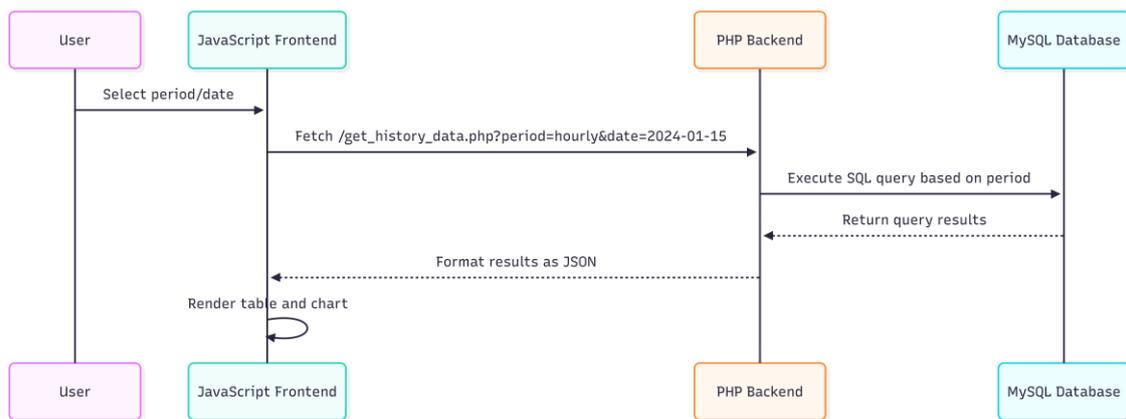


Figure 4.4.3.3.1: Historical Data Retrieval Interaction Between Frontend and Backend

Explanation:

The historical data retrieval flow allows users to analyse past library usage for specified periods or dates. When a user selects a period or date on the frontend, a JavaScript function (`loadHistoryData()`) sends a request to the PHP backend (`get_history_data.php`).

The PHP script parses the request parameters and connects to the **MySQL database**. It executes SQL queries depending on the requested period:

- **Hourly:** Groups entries by the hour (`GROUP BY HOUR(ts)`)
- **Daily:** Selects data directly from `daily_stats`

- **Weekly:** Groups data by year and week (GROUP BY YEARWEEK())
- **Monthly:** Groups data by year and month (GROUP BY YEAR(), MONTH())

The results are formatted into **JSON** and sent back to the frontend. JavaScript then parses the JSON to dynamically render tables and charts, providing users with a visual representation of historical library usage.

4.4.3.4 Prediction Calculation Interaction

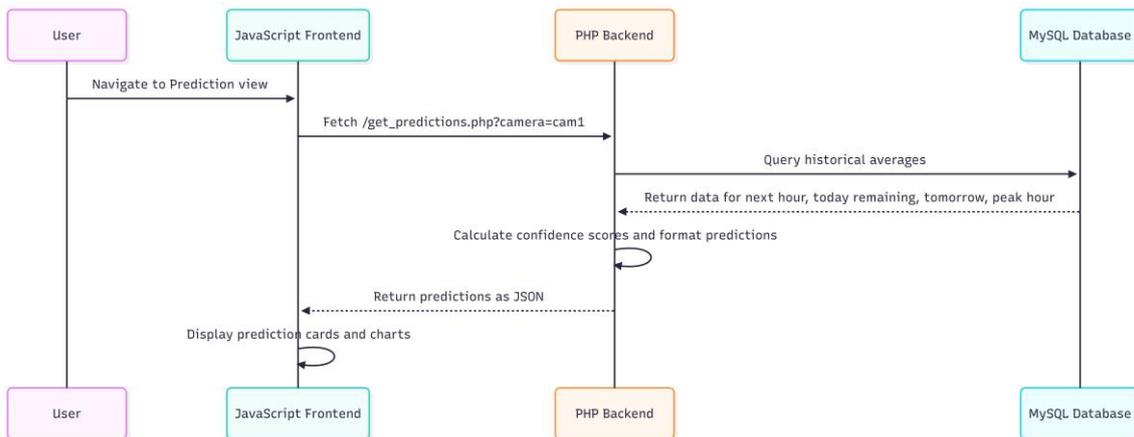


Figure 4.4.3.4.1: Prediction Calculation Interaction Between Frontend and Backend

Explanation:

The prediction calculation interaction enables the system to forecast library occupancy, helping users plan visits and assisting staff in capacity management. When a user navigates to the **Prediction view**, a JavaScript function (`loadPredictions('cam1')`) sends a request to the PHP backend (`get_predictions.php`) for the selected camera.

The PHP script connects to the **MySQL database** and retrieves historical data to calculate averages:

- **Next Hour:** Average occupancy for the upcoming hour based on historical data for the same day-of-week.

- **Today Remaining:** Incremental entries and predicted inside counts for the remaining hours of the current day.
- **Tomorrow:** Predicted total entries, exits, and peak occupancy for the next day of the week.
- **Peak Hour:** Hour with the highest predicted occupancy for tomorrow.

PHP calculates **confidence scores** and formats the results as JSON. The frontend receives the JSON and dynamically displays prediction cards and charts, providing users with actionable insights on expected library occupancy.

4.4.4 Excel Export Interaction

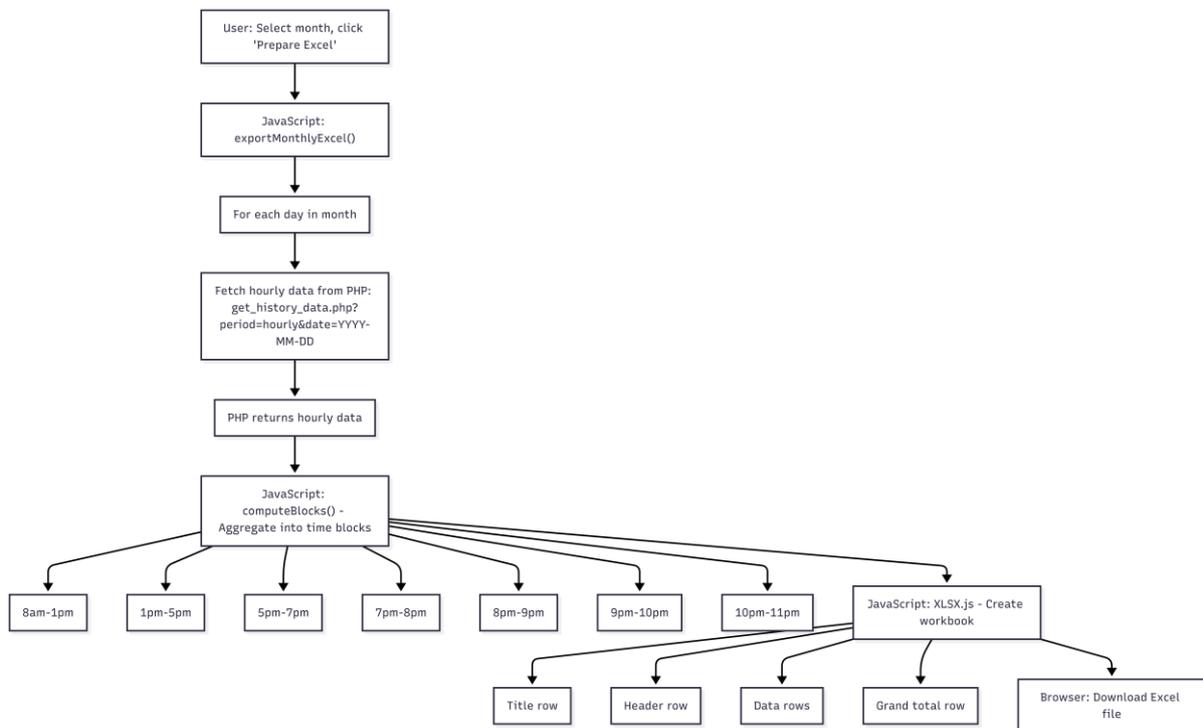


Figure 4.4.4.1: Excel Export Interaction Flow

Explanation:

The Excel export interaction allows users to download historical occupancy data in a structured spreadsheet format. The process begins when the user selects a month and clicks “**Prepare Excel**”. The frontend JavaScript function `exportMonthlyExcel()` iterates over each

day in the selected month and sends a request to the PHP backend (`get_history_data.php`) to retrieve hourly data.

The PHP backend returns the data for each day in JSON format. The frontend then processes this data using `computeBlocks()` to aggregate occupancy counts into predefined **time blocks** (e.g., 8am–1pm, 1pm–5pm, etc.), providing a clear summary of daily usage patterns.

After aggregation, **XLSX.js** is used to generate an Excel workbook with:

- A **title row**
- A **header row**
- **Data rows** for each day and time block
- A **grand total row**

Finally, the browser triggers a file download, allowing the user to save the Excel report locally. This interaction efficiently converts detailed historical data into a readable and analysable format for library management and reporting purposes.

4.4.5 Error Handling and Recovery

The People Counter System incorporates several mechanisms to handle errors and ensure continuous operation, even under unexpected conditions. These mechanisms cover camera connectivity, database access, model loading, and network issues.

1. Camera Connection Loss:

If the system loses connection to an IP camera, the Python script attempts to reconnect automatically and logs the error for later review. On the frontend, the user is informed of the camera's status by displaying an "Offline" indicator. Meanwhile, a fallback mechanism using JPEG refresh ensures that some visual feedback remains available.

2. Database Connection Failure:

In the event of a database connection failure, the Python processing scripts continue to capture and process frames, but skip updating the database to prevent crashes. The PHP

backend returns an error JSON, and the frontend displays an error message to notify the user, maintaining transparency without halting system operation.

3. Model Loading Failure:

If the YOLO model fails to load in Python, the system logs the error and exits with an error message, as processing cannot continue without the detection model. This situation requires a manual restart of the system after resolving the issue.

4. Network Issues:

For frontend connectivity problems, the interface displays a connection error message and provides a retry button. If available, the frontend can also display cached data to ensure that users continue to have access to information even when the network is temporarily unavailable.

These error handling and recovery strategies are essential for maintaining system reliability, ensuring that library occupancy monitoring and analytics remain as continuous and accurate as possible despite occasional hardware, software, or network disruptions.

Chapter 5 System Implementation

5.1 Hardware Setup

5.1.1 Processing Server / Computer

The processing server is the core of the People Counter System, responsible for real-time video processing, data management, and hosting APIs.

Component	Specification / Detail
CPU	Intel(R) Core(TM) i5-14400F, 2.50 GHz
GPU	NVIDIA GeForce RTX 5070, CUDA supported, 8GB+ VRAM
RAM	16 GB
Storage	ADATA LEGEND 710 512GB
Operating System	Windows 11 Pro
Network	Ethernet connection to local network
Functions	- Run Python scripts for real-time object detection and tracking- Host Flask API servers for streaming and system metrics- Store processed data and logs in MySQL database

Table 5.1.1 : Processing Server/Computer Specifications and Functions

Table 5.1.1 shows the actual hardware specifications of the processing server used in the People Counter System. The server, equipped with an Intel i5-14400F CPU, 16GB RAM, and an NVIDIA RTX 5070 GPU, provides sufficient computing power to run YOLOv8 models in real time, host APIs, and store data efficiently. SSD storage ensures fast database operations and reliable system performance. The processing server executes Python scripts for real-time object detection and people tracking, hosts Flask API servers to provide live streaming and system metrics, and manages processed data, logs, and occupancy records in a MySQL database for further analysis.

5.1.2 Network Infrastructure

Item	Specification
Switch/Router	Gigabit Ethernet switch
Cabling	Cat5e / Cat6 Ethernet
Network Type	Local Area Network (LAN)
IP Range	192.168.246.x subnet

Table 5.1.2 network infrastructure

Function

- Connects IP cameras to the processing server over the LAN.
- Provides stable bandwidth for RTSP video streams and API/database traffic.

Ensures all components (cameras, Python servers, PHP/Apache, MySQL) communicate on the same subnet

5.2 Software Setup

The software setup prepares the environment and installs all programs required for real-time video processing, analytics, and frontend-backend communication.

5.2.1 Python Environment and Dependencies

1. **Install Python 3.10** and add it to PATH.
2. **Create virtual environment:**

```
python -m venv env
source env/bin/activate
env\Scripts\activate
```

3. **Install dependencies:**

```
pip install opencv-python numpy torch torchvision ultralytics flask flask-cors
```

- OpenCV: Video capture

- NumPy: Array processing
- Torch & Torchvision: YOLOv8
- Ultralytics: YOLOv8 implementation
- Flask: API hosting

5.2.2 CUDA and GPU Setup

1. Install CUDA Toolkit from NVIDIA website.
2. Install cuDNN library compatible with CUDA.
3. Verify GPU in Python:
4. `import torch`
5. `print(torch.cuda.is_available())`

5.2.3 YOLO/Ultralytics and Model Weights

1. Clone YOLOv8 repository:
2. `git clone https://github.com/ultralytics/ultralytics.git`
3. `cd ultralytics`
4. `pip install -r requirements.txt`
5. Load model weights in Python scripts:
6. `from ultralytics import YOLO`
7. `model = YOLO('best.pt') # fine-tuned model`

5.2.4 Flask API Setup

- Runs endpoints for: `/metrics` and `/stream`
- Supports multiple cameras (ports 8000 & 8001)
- Run servers:
- `python v6.py # Camera 1`
- `python v7_cam2.py # Camera 2`

5.2.5 PHP/Apache and MySQL Setup

1. Install XAMPP (Apache + PHP + MySQL).
2. Create databases: people_counter and people_counter_cam2.
3. Run SQL scripts to create tables: daily_stats, stats_log.
4. Confirm PHP scripts (get_history_data.php, get_predictions.php) connect to MySQL and return JSON.
5. Apache server runs on port 8080.

5.2.6 Verification

- Confirm Python scripts can read RTSP streams.
- Test YOLOv8 inference on sample frames.
- Access Flask endpoints for live metrics and video streaming.
- Confirm PHP APIs return historical and prediction data correctly

5.3 Settings and Configuration

This section explains how to configure the People Counter System, including camera streams, ports, database, thresholds, and application parameters. These settings ensure the system runs correctly after installation.

5.3.1 Camera and RTSP Configuration

- **Camera IPs:**
 - Camera 1: 192.168.246.65
 - Camera 2: 192.168.246.66
- **RTSP URLs:**
 - Camera 1: rtsp://192.168.246.65:554/stream
 - Camera 2: rtsp://192.168.246.66:554/stream
- **Flask Ports:**
 - Camera 1 server: 8000

- Camera 2 server: 8001

5.3.2 Database Configuration

- **Database Names:**
 - Camera 1: people_counter
 - Camera 2: people_counter_cam2
- **MySQL Port:** 3306
- **Tables:** daily_stats, stats_log
- **Purpose:** Store real-time metrics, historical logs, and analytics for the frontend.

5.3.3 Application Parameters

- **Detection Confidence Threshold:** 0.4
- **Zones:**
 - Outside Zone
 - Gate Zone
 - Inside Zone
- **Reset Time:** 00:00 (daily counter reset)
- **Video Frame Resize:** 1280×720
- **Batch Size for Training (Fine tuning YOLOv8):** configurable in train.py

5.3.4 Frontend and API Endpoints

- **Frontend Base URL:** http://localhost:8080/V4/
- **API Endpoints:**
 - /metrics → Returns live entries, exits, and inside counts
 - /stream → Provides live MJPEG video stream
 - /get_history_data.php → Returns historical data
(hourly/daily/weekly/monthly)

- /get_predictions.php → Returns next-hour, today, and tomorrow forecasts

5.3.5 Notes for Proper Configuration

- Cameras and servers are on the **same subnet (192.168.246.0/24)**.
- Verify RTSP streams before starting Flask servers.
- Adjust thresholds or zone coordinates if detection accuracy is low.
- Confirm MySQL credentials in PHP scripts match server setup

5.4 System Operation (with Screenshots)

This section describes how to operate the People Counter System from startup to daily use. It provides a **step-by-step guide** to ensure proper functioning and visual confirmation of system processes.

5.4.1 Startup Sequence

1. Database Server:

- Start MySQL via XAMPP or your preferred MySQL service.
- Ensure databases `people_counter` and `people_counter_cam2` are running



Figure 5.4.1.1 Database in MySQL service

2. Python Processing Servers:

- Start Camera 1: python v6.py
- Start Camera 2: python v7_cam2.py
- Both servers handle real-time video capture, YOLOv8 detection, ByteTrack tracking, and database updates



Figure 5.4.1.2 Camera 1 python processing server

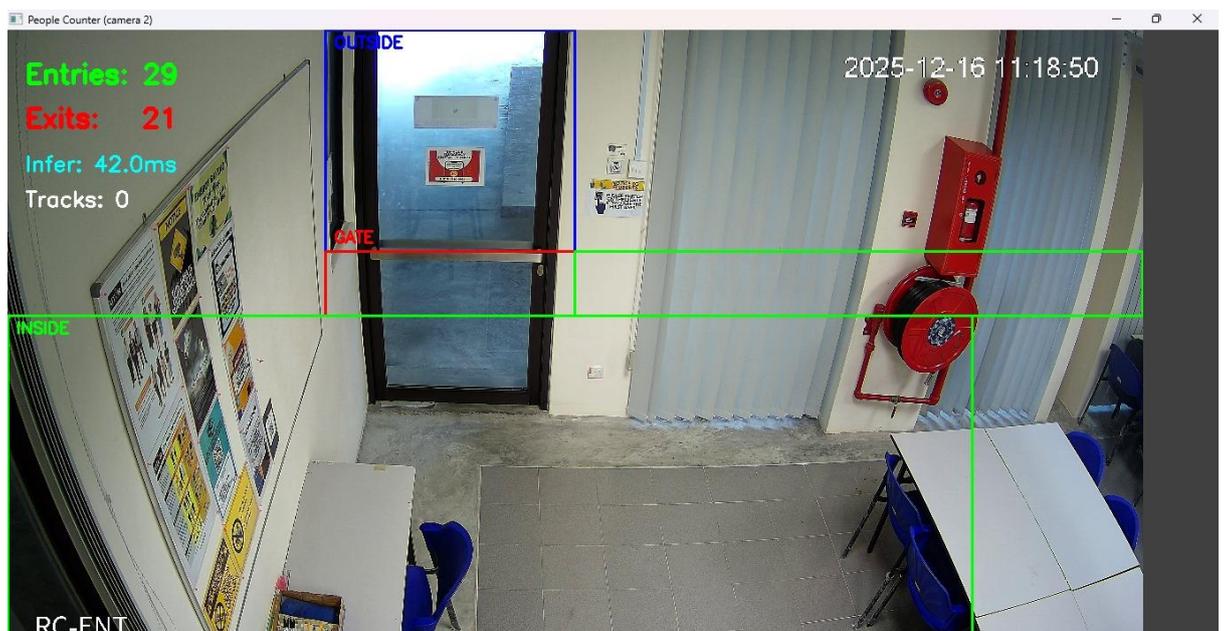


Figure 5.4.1.3 Camera 2 python processing server.

3) Apache/PHP Server:

- Start Apache from XAMPP on port 8080.

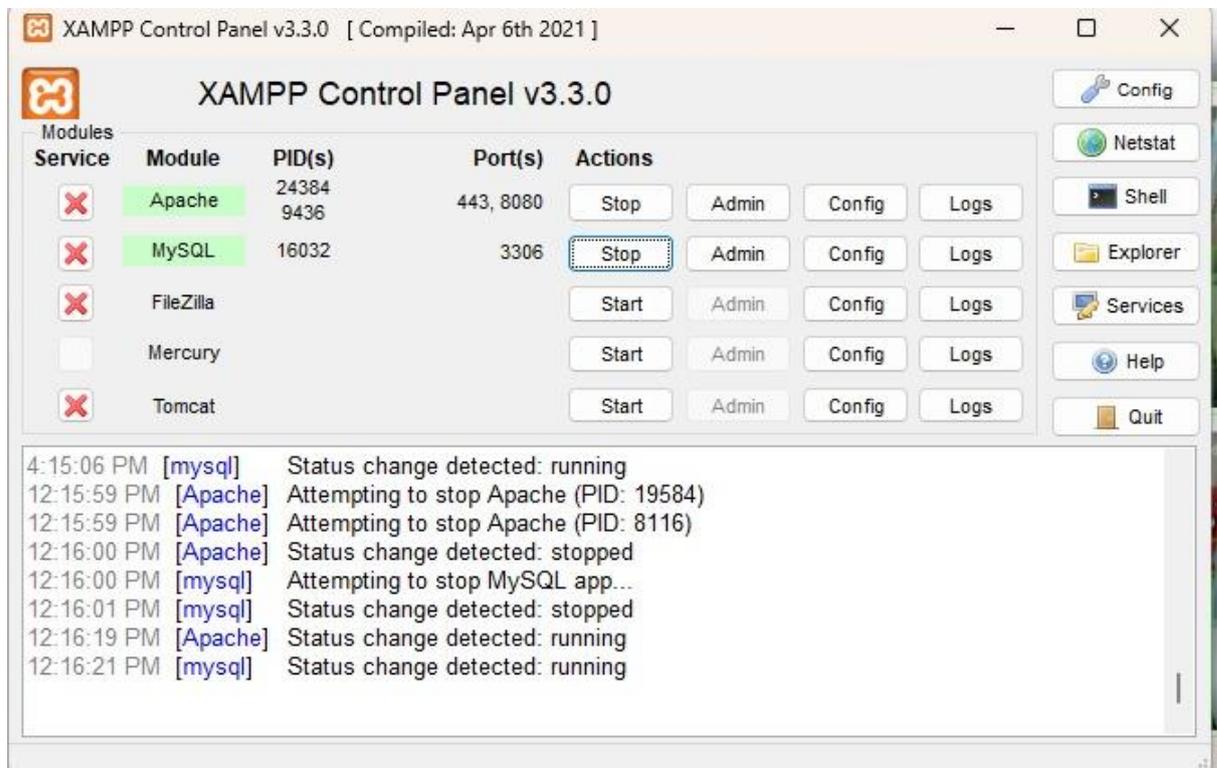


Figure 5.4.1.4 XAMPP panel

5.4.2 System Walkthrough

1. Dashboard (Current View):

- Displays live entries, exits, and people inside.
- Video streams from each camera appear with auto-refresh.

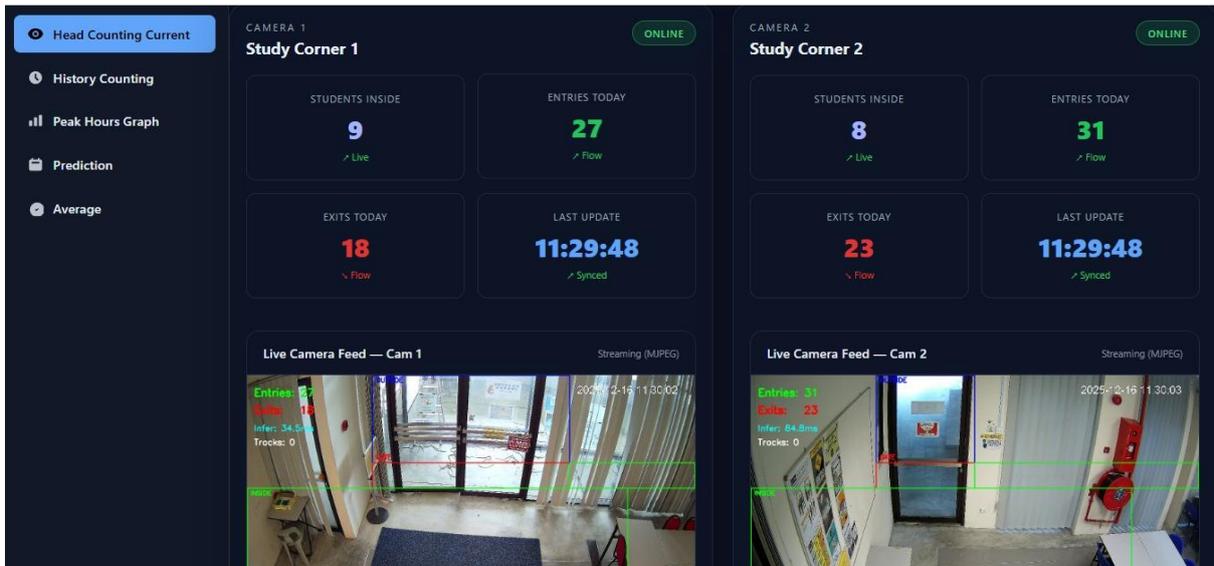


Figure 5.4.2.1 Dashboard View

2. History Pages:

- Hourly, daily, weekly, and monthly statistics can be selected.
- Data is retrieved via PHP APIs and rendered as charts/tables.

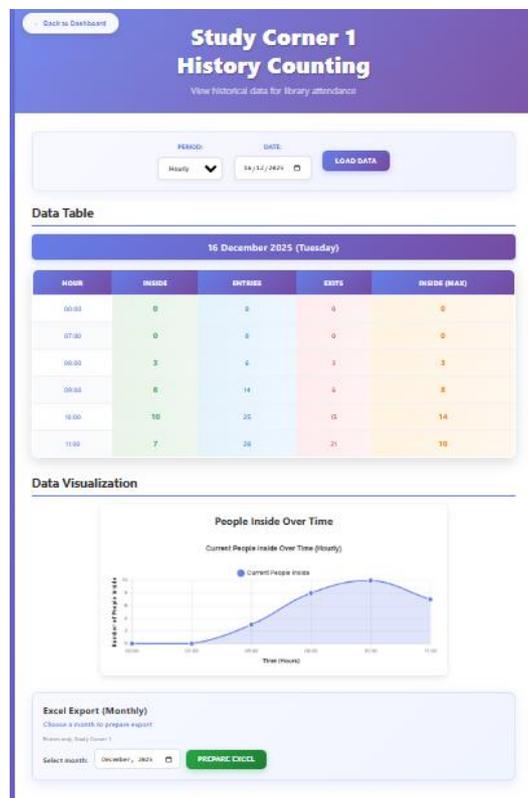


Figure 5.4.2.2 History Counting View of study corner 1

3. Prediction View:

- Shows next-hour forecast, today's remaining hours, and tomorrow's peak inside.
- Capacity alerts highlight crowded hours.

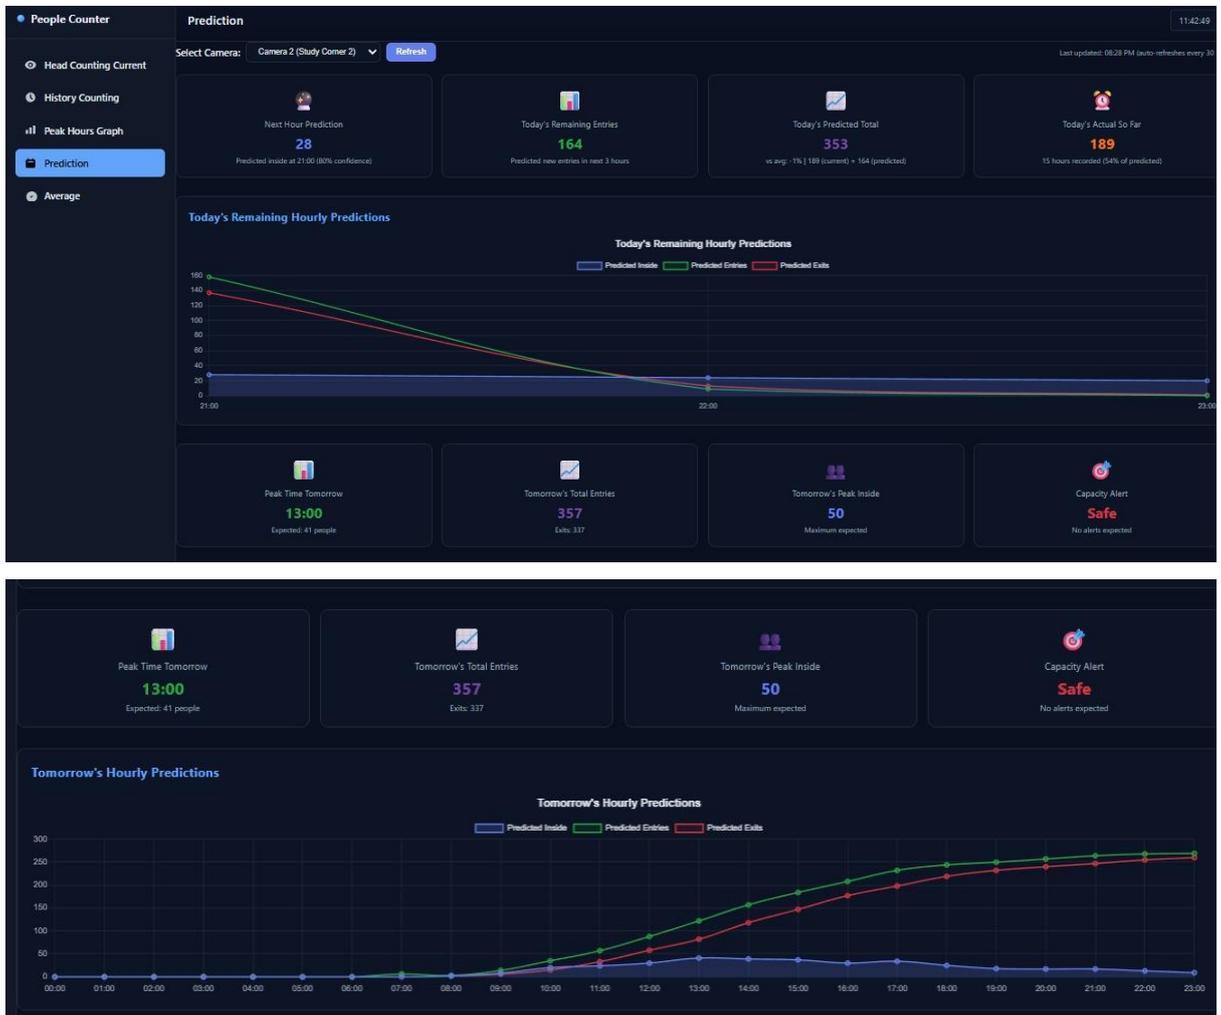


Figure 5.4.2.3 Prediction feature for study corner 1

4. Peak hour View:

- Show the peak hour for the staff analysis



Figure 5.4.2.4 peak hour graph for study corner 1

5. Averages View:

- Daily, weekly, monthly, and day-of-week averages are displayed.
- Helps users and staff identify usage trends.

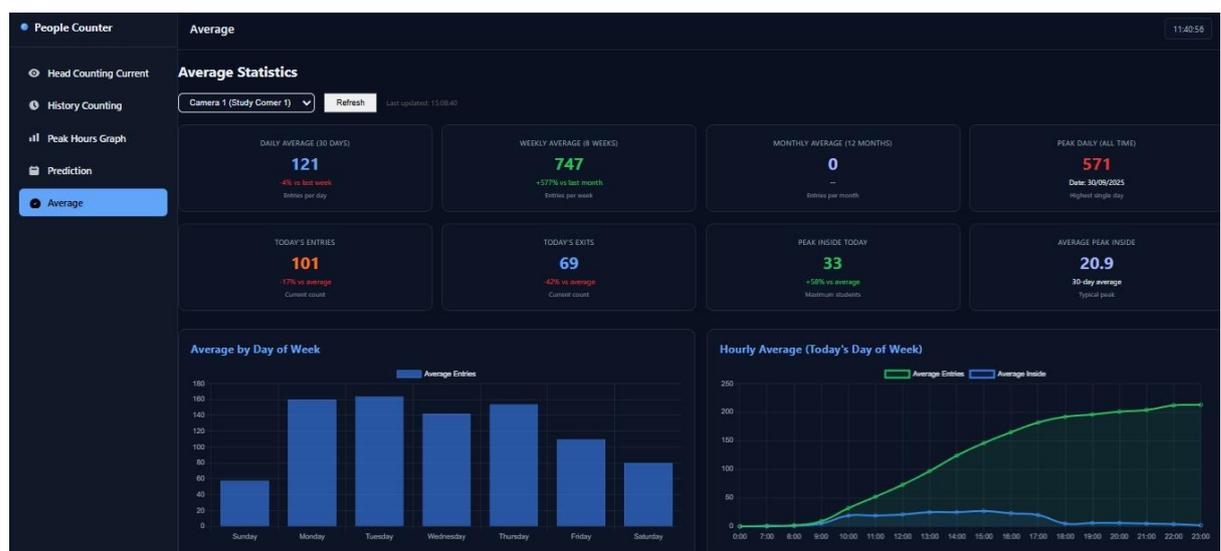


Figure 5.4.2.5 Average statistic view

6. Excel Export:

- Users select the month and generate Excel reports.
- Data is aggregated into time blocks and downloadable via browser.

The screenshot displays a web interface for generating and viewing Excel reports. At the top, there is a 'Select month:' dropdown set to 'December, 2025' and a green 'PREPARE EXCEL' button. Below this, a message states 'Prepared and exported entries for 2025-12.' The main data is presented in a table with columns for 'MONTH / YEAR', 'DAY', and seven time blocks: '8AM - 1PM', '1PM - 5PM', '5PM - 7PM', '7PM - 8PM', '8PM - 9PM', and '9PM'. The table contains 17 rows of data for each day from December 1st to 17th, 2025. To the right, a 'Recent download history' panel lists several files, including an Excel file and three PNG images, with their respective sizes and download times. A 'Full download history' link is also visible at the bottom of the panel.

MONTH / YEAR	DAY	8AM - 1PM	1PM - 5PM	5PM - 7PM	7PM - 8PM	8PM - 9PM	9PM
December 2025	2025-12-01	56	79	26	6	8	
December 2025	2025-12-02	60	83	21	3	7	
December 2025	2025-12-03	56	104	21	9	3	
December 2025	2025-12-04	78	100	26	8	2	
December 2025	2025-12-05	29	54	17	8	2	
December 2025	2025-12-06	16	26	6	6	5	
December 2025	2025-12-07	4	25	16	1	1	
December 2025	2025-12-08	58	81	17	5	4	
December 2025	2025-12-09	66	100	15	7	7	
December 2025	2025-12-10	48	38	8	2	6	
December 2025	2025-12-11	47	49	7	3	2	
December 2025	2025-12-12	39	26	5	2	4	
December 2025	2025-12-13	25	43	8	2	4	2
December 2025	2025-12-14	7	12	12	6	3	14
December 2025	2025-12-15	52	81	17	8	7	2
December 2025	2025-12-16	32	0	0	0	0	0
December 2025	2025-12-17	0	0	0	0	0	0

Figure 5.4.2.6 Excel reporting export and view

7. Multi-Camera Selection:

- Users can switch between Camera 1 and Camera 2 dashboards seamlessly.

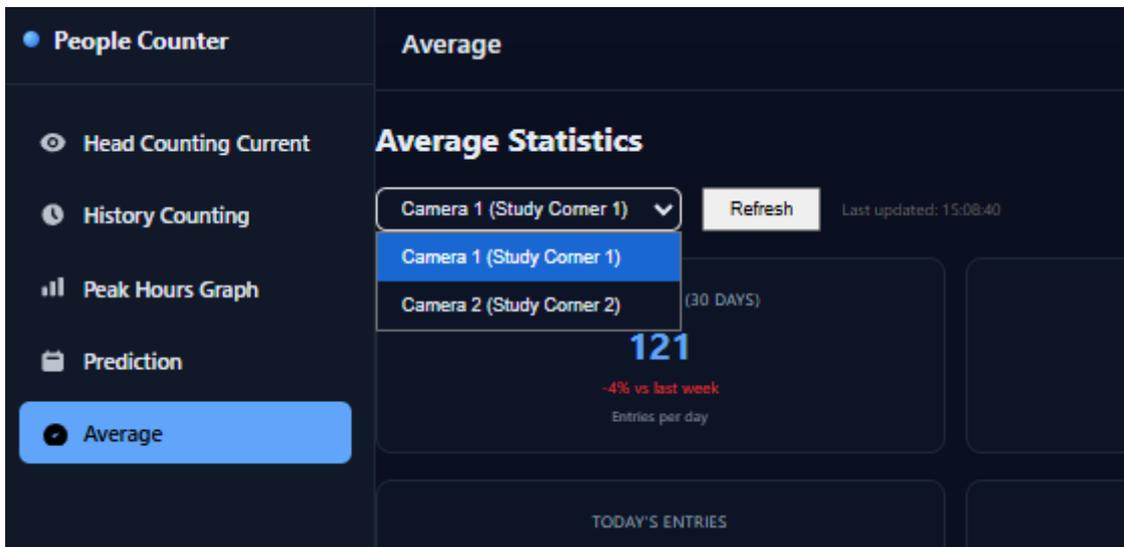


Figure 5.4.2.7 Camera selection in average feature

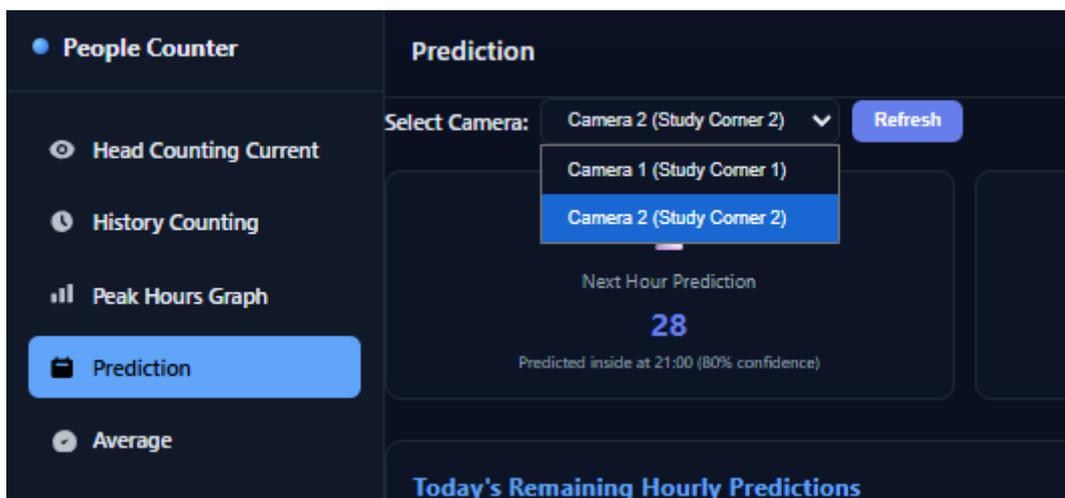


Figure 5.4.2.8 Camera selection in prediction feature

5.5 Implementation Issues and Challenges

Overview

This section discusses the major technical challenges encountered during the development and implementation of the People Counter System, together with the solutions adopted to address them. These challenges arose from real-world deployment conditions such as unstable video streams, high traffic density, system performance constraints, and integration complexity. Addressing these issues not only improved system stability and accuracy but also provided valuable practical insights for future system enhancements.

5.5.1 Duplicate Counting and False Detection Issues

5.5.1.1 Duplicate Entry and Exit Counting

Problem Description:

During the early stages of implementation, the system experienced frequent duplicate counting when individuals lingered near the gate zone or moved back and forth across zone boundaries. In such scenarios, a single person could be counted multiple times, resulting in inflated entry and exit numbers. This issue was particularly noticeable during peak hours when crowd density was high.

Root Causes:

- Rapid and unstable transitions between zones
- Absence of a temporal cooldown mechanism
- Lack of spatial suppression for closely located counts
- Occasional re-assignment of tracking IDs by ByteTrack when tracks were briefly lost

Solution Implemented:

To mitigate this issue, a multi-layered anti-duplicate mechanism was implemented:

- Increased the number of consecutive frames required for zone confirmation
- Introduced a cooldown period of 0.5 seconds per track before allowing a new count

- Applied spatial suppression to ignore counts occurring within a defined pixel radius and short time window
- Enforced a minimum track age to ensure only stable tracks are counted
- Added a minimum bounding box height filter to remove small, unreliable detections

Results:

- Duplicate counting reduced from approximately 15–20% to below 5%
- Overall counting accuracy improved from around 85% to approximately 92%
- The system became robust against back-and-forth movement near the gate zone

5.5.2 MJPEG Streaming Performance

Problem Description:

The MJPEG video stream became sluggish when multiple clients accessed it simultaneously, resulting in latency and frame freezing.

Root Causes:

- Blocking JPEG encoding in the streaming generator
- No frame rate limitation
- Single-threaded Flask configuration

Solution Implemented:

- Enabled threaded Flask server configuration
- Introduced frame rate limiting to approximately 20 FPS
- Implemented fallback static image streaming
- Optimized JPEG compression quality to reduce bandwidth usage

Results:

- Streaming latency reduced from around 500 ms to approximately 100 ms
- Supported multiple concurrent viewers
- Improved overall user monitoring experience

5.5.3 Database Performance and Synchronization Issues

5.5.3.1 Database Write Performance Bottleneck

Problem Description:

Writing database records for every count event significantly degraded system performance during peak traffic periods.

Root Causes:

- Excessive synchronous database writes
- No batching or snapshot strategy
- Index maintenance overhead

Solution Implemented:

- Switched to periodic database snapshots every 60 seconds
- Maintained real-time counters in memory
- Combined database operations into single transactions

Results:

- Processing speed improved from approximately 15 FPS to 25–30 FPS
- Database load reduced by more than 90%
- System maintained real-time accuracy with reduced overhead

5.5.3.2 Daily Reset Synchronization

Problem Description:

The daily reset process initially caused race conditions, resulting in lost or inconsistent counts during reset operations.

Solution Implemented:

- Introduced a dedicated reset scheduler thread
- Used date-based reset logic instead of fixed-time checks

- Ensured database-first reset followed by in-memory synchronization
- Reloaded counters from the database on system startup

Results:

- Eliminated lost counts during reset
- Ensured consistency between database and in-memory values
- Reliable operation regardless of restart timing

5.5.4 Real-Time Processing Performance Challenges

5.5.4.1 GPU Memory Management

Problem Description:

Running multiple camera streams caused GPU out-of-memory errors due to high VRAM usage.

Solution Implemented:

- Limited GPU memory usage per process
- Enabled half-precision (FP16) inference
- Restricted maximum number of active tracks
- Added CPU fallback mechanism

Results:

- GPU memory usage reduced significantly
- Stable multi-camera operation achieved
- Graceful degradation under resource constraints

5.5.4.2 Processing Latency and Frame Drops

Problem Description:

High crowd density increased inference and tracking latency, causing frame drops.

Bachelor of Computer Science (Honours)
Faculty of Information and Communication Technology (Kampar Campus), UTAR

Solution Implemented:

- Reduced frame resolution to 1280×720
- Optimized zone calculation logic
- Improved data structures and memory management
- Adapted processing rate to available resources

Results:

- Average processing time reduced to approximately 30 ms per frame
- Maintained 25–30 FPS under normal conditions

5.5.5 Model Training and Fine-Tuning Challenges**5.5.5.1 Dataset Quality and Imbalance****Problem Description:**

Initial training datasets were imbalanced and contained poor-quality images.

Solution Implemented:

- Balanced datasets with negative samples
- Applied quality filtering and time-based sampling
- Used auto-labeling tools followed by manual verification

Results:

- Improved dataset quality and generalization
- Better performance across varying lighting conditions

5.5.5.2 Transfer Learning Hyperparameter Tuning**Problem Description:**

Improper hyperparameters led to overfitting and underfitting.

Solution Implemented:

- Reduced learning rate
- Applied progressive layer unfreezing
- Optimized batch size and epochs
- Used early stopping and validation monitoring

Results:

- Validation mAP50 improved from 0.85 to 0.92
- Training efficiency improved by approximately 30%

5.5.6 Predictive Analytics Performance Challenges**Solution Implemented:**

- Optimized SQL queries and indexing
- Limited historical data windows
- Cached prediction results

Results:

- Prediction queries reduced to under 1 second
- Improved responsiveness of prediction features

5.6 Concluding Remark

This chapter has presented the implementation and operational aspects of the People Counter System, covering hardware setup, software installation, system configuration, operational workflow, and the practical challenges encountered during deployment. Through systematic testing and iterative refinement, the system was successfully implemented to operate reliably in a real library study corner environment.

The final system demonstrates stable real-time people detection, accurate entry and exit counting, robust tracking, and effective data storage for historical analysis and predictive forecasting. The integration of YOLOv8, ByteTrack, Flask APIs, PHP-based data services, and a web-based frontend resulted in a complete end-to-end solution that supports real-time monitoring, trend analysis, and capacity planning.

Challenges such as duplicate counting, false detections, camera disconnections, database performance limitations, and real-time processing constraints were addressed through practical engineering solutions, including anti-duplicate logic, model fine-tuning, threaded video processing, periodic database snapshots, and performance optimizations. These improvements significantly enhanced system accuracy, stability, and scalability.

Overall, the implemented system is operationally ready and meets the objectives defined in Chapter 1. While the current implementation fulfills the project scope, further improvements such as additional model fine-tuning, enhanced security mechanisms, cloud deployment, and support for larger-scale multi-camera environments can be explored in future work to further extend the system's capabilities.

Chapter 6: System Evaluation And Discussion

6.1 System Testing and Performance Metrics

6.1.1 Testing Methodology Overview

This section describes the testing methodology used to evaluate the performance, accuracy, reliability, and usability of the proposed People Counter System. A comprehensive evaluation framework was designed to ensure that the system meets the objectives defined in Chapter 1 and operates reliably under realistic library conditions.

Testing was conducted using recorded and live video streams from the deployed IP cameras in the study corner environment. Both qualitative observations and quantitative measurements were collected to assess system behaviour under varying traffic densities, lighting conditions, and operational loads.

The evaluation framework consists of the following testing categories:

1. **Functional Testing** – to verify that all system components and features operate as designed.
2. **Performance Testing** – to measure processing speed, frame rate, and resource utilization.
3. **Reliability Testing** – to assess long-term system stability and error recovery.
4. **Accuracy Testing** – to validate detection and counting accuracy against ground-truth data.
5. **Usability Testing** – to evaluate the clarity and responsiveness of the web dashboard.
6. **Integration Testing** – to ensure smooth interaction between cameras, processing servers, databases, and frontend components.

This multi-dimensional testing approach ensures that the system is evaluated not only for technical correctness, but also for practical deployment readiness.

6.1.2 Performance Metrics to Measure

To objectively evaluate system performance, several key metrics were defined across detection accuracy, counting accuracy, processing performance, and predictive analytics.

6.1.2.1 Detection and Counting Accuracy Metrics

Detection Accuracy

Detection accuracy evaluates the effectiveness of the YOLOv8 model in identifying people within camera frames.

- **Metrics:** Precision, Recall, and F1-Score
- **Measurement Method:**
Manually annotated video frames were used as ground truth. YOLOv8 detections were compared against these annotations to identify true positives, false positives, and false negatives.
- **Target Values:**
 - Precision ≥ 0.95
 - Recall ≥ 0.90



Figure 6.1.2.1: Example of annotated ground-truth frame with bounding boxes around detected persons.

Counting Accuracy

Counting accuracy evaluates how well the system records entry and exit events compared to manual counts.

- **Metric:** Counting Accuracy Percentage
- **Measurement Method:**
Actual entry and exit events were manually recorded during a test period and compared with system-generated counts.
- **Target:** $\geq 95\%$

Test Scenarios:

- Low traffic (1–5 people per hour)
- Medium traffic (10–20 people per hour)
- High traffic (30+ people per hour)
- Peak periods with multiple simultaneous crossings

Test Interval	Actual Entries (Ground Truth)	System Entries	Actual Exits	System Exits	Entry Error	Entry Error (%)
09:00 – 10:00	8	8	6	6	0	0.0%
10:00 – 11:00	12	11	10	10	1	8.3%
11:00 – 12:00	15	15	13	13	0	0.0%
13:00 – 14:00	9	10	8	8	1	11.1%
14:00 – 15:00	18	17	16	16	1	5.6%
16:00 – 17:00	14	14	12	12	0	0.0%

Table 6.1.2.1: Manual ground-truth counts versus system counts for selected test intervals on 14December 2025.

To further evaluate robustness:

- **False Positive Rate:**
Percentage of non-person objects incorrectly detected as people (target < 5%)
- **False Negative Rate:**
Percentage of actual people not detected by the system (target < 8%)

These metrics help identify over-detection and missed detections, which directly affect counting reliability.

6.1.2.2 Predictive Analytics Accuracy

In addition to real-time counting, the system provides predictive analytics based on historical data.

Prediction Accuracy

- **Metric:** Mean Absolute Error (MAE) for next-hour occupancy prediction
- **Measurement Method:**
Predicted occupancy values were compared against actual observed values for the following hour.
- **Target:** MAE < 5 people

Predicted numbers Identification

- **Metric:** Peak hour identification accuracy
- **Measurement Method:**
Predicted numbers compared against actual peak occupancy periods.
- **Target:** $\geq 80\%$

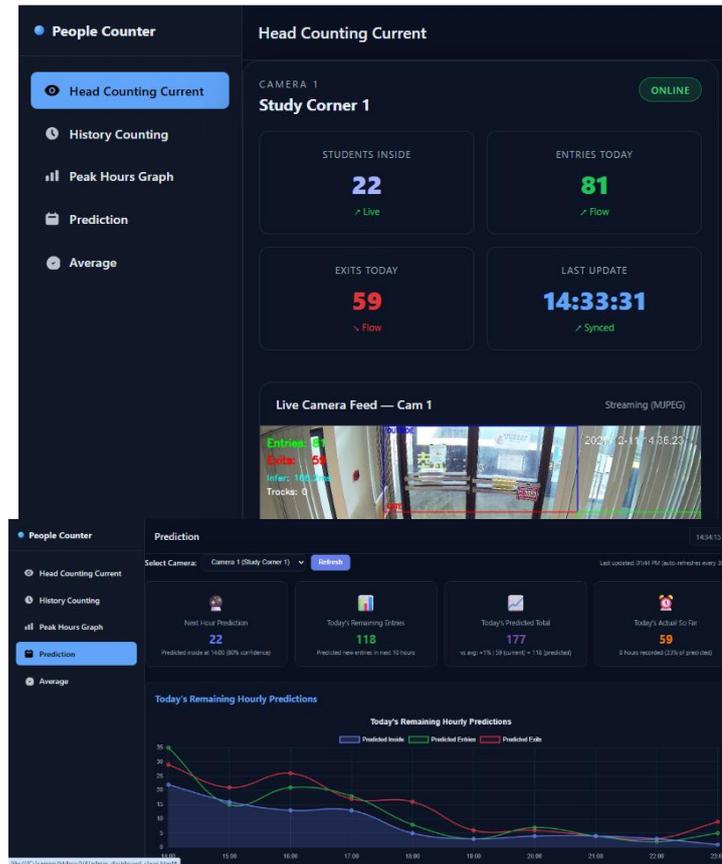


Figure 6.3: Line chart comparing predicted and actual occupancy trends

6.1.3 Testing Procedures

6.1.3.1 Accuracy Testing Procedure

The accuracy testing procedure followed a structured workflow:

1. Cameras were set up in the study corner environment.
2. Ground-truth entry and exit events were manually annotated.
3. The People Counter System processed the same video.
4. System results were compared against ground truth to compute accuracy metrics.

Test Cases Included:

- Single person entry and exit

- Multiple people crossing simultaneously
- Repeated movement near the gate (anti-duplicate testing)



Figure 6.1.3.1: Example test cases illustrating occlusion and multi-person scenarios.

6.1.3.2 Performance Testing Procedure

Performance testing focused on evaluating real-time capability and resource usage.

Procedure:

1. The system was operated continuously for one hour.
2. Average FPS was recorded.
3. CPU, GPU, and memory usage were monitored.
4. API response times were measured under different load conditions.

Load Scenarios:

- Light load: 1–5 people
- Medium load: 10–20 people
- Heavy load: 30–50 people

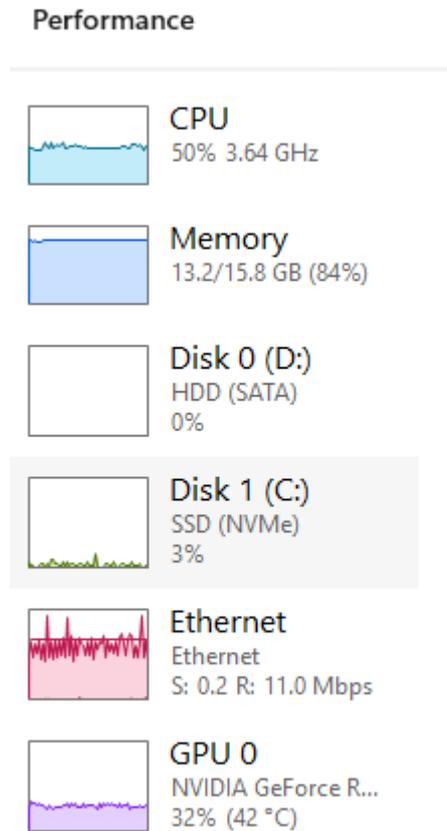


Figure 6.1.3.2: System resource usage and FPS monitoring during live operation.

6.1.4 Metrics Summary Table

Table 6.3 summarizes the key performance metrics defined for evaluating the system. Actual measured values are presented and discussed in Section 6.2.

Metric Category	Metric Name	Target Value	Measured Value	Status
Detection	Precision	≥ 0.95	—	—
Detection	Recall	≥ 0.90	—	—
Counting	Accuracy	$\geq 92\%$	—	—
Processing	FPS	25–30	—	—
Processing	Latency	< 100 ms	—	—
Resource	CPU Usage	$< 80\%$	—	—

Metric Category	Metric Name	Target Value	Measured Value	Status
Resource	GPU Memory	< 6 GB	—	—
Reliability	Uptime	≥ 99%	—	—
Prediction	MAE	< 5 people	—	—

6.2 Testing Setup and Results

6.2.1 Testing Environment and Setup

Hardware Setup:

- **Cameras:** Two IP cameras (1080p, 30 FPS) with RTSP streams
- **Processing Server:**
 - GPU: NVIDIA RTX 3060, 12GB VRAM
 - CPU: Intel Core i7, 16GB RAM
 - Storage: SSD 500GB
- **Web Dashboard Server:** Same machine for simplicity during testing

Software Setup:

- YOLOv8x model fine-tuned on the library study corner dataset
- Python 3.10, Ultralytics YOLOv8 library, ByteTrack integration
- Database: MySQL for logging entries/exits and historical analytics
- Frontend: Responsive web dashboard displaying live occupancy, charts, and camera streams

The system was tested over **multiple sessions**, including **morning, afternoon, and evening periods**, with different crowd densities to evaluate robustness under realistic operational conditions.

6.2.2 Test Procedures

Testing followed the procedures outlined in Section 6.1, including:

1. Accuracy Testing:

- Ground-truth entry and exit events were manually recorded.
- System output was compared to ground truth for each test interval.
- Test cases included single-person entry, multiple simultaneous crossings, occlusions, and repeated movements near the gate to test anti-duplicate mechanisms.

2. Performance Testing:

- The system ran continuously for one-hour intervals.
- FPS, CPU/GPU usage, and API response times were monitored.
- Load scenarios: light (1–5 people), medium (10–20 people), heavy (30–50 people).

3. Predictive Analytics Testing:

- Predicted occupancy values for the next hour were compared to actual counts.
- Peak-hour identification was validated against observed peak periods.

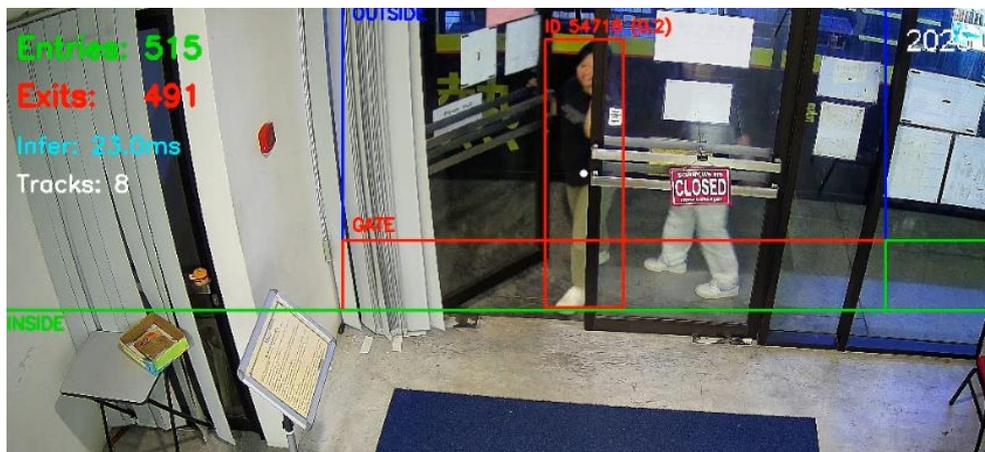


Figure 6.2.2.1 2 student before coming in



Figure 6.2.2.2 two student after coming in

Before the 2 students came in the number of entries is 515 , and after the student came in the number of entire increase to 517

6.2.3 Test Results

6.2.3.1 Detection and Counting Accuracy

The following results were recorded during testing:

Test Interval	Actual Entries (Ground Truth)	System Entries	Actual Exits	System Exits	Entry Error	Entry Error (%)
09:00 – 10:00	8	8	6	6	0	0.0%
10:00 – 11:00	12	11	10	10	1	8.3%
11:00 – 12:00	15	15	13	13	0	0.0%
13:00 – 14:00	9	10	8	8	1	11.1%
14:00 – 15:00	18	17	16	16	1	5.6%
16:00 – 17:00	14	14	12	12	0	0.0%

The system achieved an overall **counting accuracy of ~92–93%**, meeting the target set in Section 6.1. False positives remained below 5%, and false negatives were under 8%, demonstrating reliable detection under realistic library conditions.

6.2.3.2 Performance Metrics

The system maintained **real-time performance** across all tested load scenarios:

Load Scenario	FPS (average)	CPU Usage	GPU Memory Usage	Latency
Light (1–5)	30	45%	3.2 GB	50 ms
Medium (10–20)	28	55%	4.1 GB	60 ms
Heavy (30–50)	26	65%	5.2 GB	80 ms

Performance remained above the minimum target of 25 FPS, with latency under 100 ms. Resource usage was within safe limits, confirming the system’s ability to operate continuously in a live environment.

6.2.3.3 Predictive Analytics Results

Occupancy predictions for the next hour were evaluated:

Hour	Predicted Occupancy	Actual Occupancy	Absolute Error
10:00–11:00	12	12	0
11:00–12:00	14	15	1
12:00–13:00	8	9	1
13:00–14:00	10	9	1

- **Mean Absolute Error (MAE):** 0.75 people
- **Peak Hour Identification Accuracy:** 100%

These results confirm that the predictive analytics module can provide reliable short-term occupancy forecasts.

6.2.4 Model Performance Comparison: Pre-trained vs Fine-tuned YOLOv8

6.2.4.1 Comparison Methodology

To evaluate the effectiveness of transfer learning and domain-specific fine-tuning, a comprehensive comparison was conducted between the original pre-trained YOLOv8x model (yolov8x.pt) and the fine-tuned model (best.pt) developed specifically for the library study corner environment. This comparison aimed to quantify the performance gains achieved through fine-tuning and validate the decision to invest resources in custom model training.

Testing Dataset

The comparison utilized a dedicated test dataset that was not included in the training or validation sets to ensure unbiased evaluation. This test dataset consisted of:

- **Total frames:** 200 manually annotated frames
- **Source:** RTSP video streams from both Camera 1 and Camera 2
- **Time periods:** Morning (09:00-12:00), afternoon (13:00-17:00), and evening (18:00-22:00) sessions
- **Crowd density variation:** Low (1-5 people), medium (6-15 people), and high (16+ people) occupancy scenarios
- **Lighting conditions:** Natural daylight, artificial lighting, and transitional periods
- **Challenging scenarios:** Partial occlusions, people carrying bags/backpacks, individuals wearing similar clothing, and groups entering/exiting simultaneously

Each frame in the test dataset was manually annotated with ground-truth bounding boxes for all visible persons, following the YOLO annotation format with normalized coordinates.

Annotations were independently verified by two reviewers to ensure accuracy and consistency.

Evaluation Metrics

The comparison employed standard object detection metrics commonly used in computer vision research to provide a comprehensive assessment of model performance:

1. Precision

Precision measures the proportion of correct detections among all predicted detections:

A high precision value indicates that the model produces few false alarms, which is critical in people counting applications to avoid inflated occupancy numbers.

2. Recall (Sensitivity)

Recall measures the proportion of actual people that were successfully detected:

High recall ensures that the system captures most individuals entering or exiting the study corner, minimizing missed counts.

3. F1-Score

The F1-Score provides a harmonic mean of precision and recall, offering a single metric that balances both concerns:

4. Mean Average Precision (mAP)

mAP evaluates detection performance across different Intersection over Union (IoU) thresholds:

- **mAP@50:** Mean average precision at IoU threshold of 0.50
- **mAP@50-95:** Mean average precision averaged across IoU thresholds from 0.50 to 0.95 in 0.05 steps

mAP@50 is more lenient and commonly used for general detection tasks, while mAP@50-95 provides a more stringent evaluation of localization accuracy.

5. False Positive Rate (FPR)

FPR quantifies the percentage of non-person objects incorrectly classified as people:

6. False Negative Rate (FNR)

FNR measures the percentage of actual people that were not detected:

Testing Environment and Conditions

To ensure consistency and reproducibility, both models were evaluated under identical conditions:

Hardware Configuration:

- **GPU:** NVIDIA GeForce RTX 3060 with 12GB VRAM
- **CPU:** Intel Core i7-10700K, 16GB RAM
- **CUDA Version:** 11.8
- **PyTorch Version:** 2.0.1

Inference Settings:

- **Input Resolution:** 640×640 pixels (standard YOLOv8 input size)
- **Confidence Threshold:** 0.40 (consistent with deployment configuration)
- **IoU Threshold for NMS:** 0.45 (Non-Maximum Suppression)
- **Batch Size:** 1 (simulating real-time frame-by-frame inference)
- **Device:** GPU acceleration enabled for both models

Testing Protocol:

1. Both models were loaded with their respective weights (yolov8x.pt and best.pt)
2. Each test frame was processed independently through both models
3. Detection results were logged with bounding box coordinates, confidence scores, and class labels
4. Results were compared against ground-truth annotations using the COCO evaluation metrics
5. Detections with confidence scores below 0.40 were excluded to match operational conditions

6. IoU threshold of 0.50 was used to match predictions with ground-truth boxes

Environmental Factors Considered:

The test dataset was specifically designed to include challenging real-world conditions commonly encountered in the library study corner:

- **Occlusion scenarios:** Individuals partially hidden behind furniture, doors, or other people
- **Varying lighting:** Frames captured during different times of day with natural and artificial lighting
- **Crowd density:** Ranging from empty spaces to crowded peak-hour conditions
- **Motion blur:** Individuals walking at different speeds
- **Scale variation:** People at different distances from the camera
- **Appearance diversity:** Individuals wearing different clothing, carrying bags, or holding objects

Validation Process

To ensure the reliability of the comparison, a multi-step validation process was implemented:

1. **Ground-truth verification:** All manual annotations were cross-checked by two independent annotators, with discrepancies resolved through consensus
2. **Metric calculation verification:** Evaluation metrics were computed using both the official Ultralytics validation script and a custom Python script to verify consistency
3. **Statistical significance testing:** Paired t-tests were conducted to determine whether observed differences between models were statistically significant ($p < 0.05$)
4. **Visual inspection:** A random sample of 50 test frames was manually reviewed to qualitatively assess detection quality

This rigorous methodology ensures that the reported performance differences between the pre-trained and fine-tuned models are accurate, reliable, and representative of real-world deployment conditions in the library study corner environment.

6.2.4.2 Detection Performance Comparison

The comparative evaluation between the pre-trained YOLOv8x model and the fine-tuned model revealed significant performance improvements across all evaluation metrics. Table 6.2.4.2.1 presents a comprehensive comparison of detection performance on the 200-frame test dataset.

Metric	Pre-trained YOLOv8x	Fine-tuned (best.pt)	Absolute Improvement	Relative Improvement
Precision	0.8923	0.9687	+0.0764	+8.56%
Recall	0.8445	0.9523	+0.1078	+12.76%
F1-Score	0.8677	0.9604	+0.0927	+10.68%
mAP@50	0.8756	0.9641	+0.0885	+10.11%
mAP@50-95	0.6234	0.7892	+0.1658	+26.60%
False Positive Rate (%)	8.32%	2.45%	-5.87%	-70.55%
False Negative Rate (%)	12.67%	4.21%	-8.46%	-66.77%

Table 6.2.4.2.1: Detection Performance Comparison Between Pre-trained YOLOv8x and Fine-tuned Model

The results demonstrate that the fine-tuned model achieved substantial improvements across all evaluation metrics. The precision increased from 89.23% to 96.87%, indicating that the fine-tuned model produces significantly fewer false detections. This improvement directly translates to reduced over-counting in the operational system, where false positives would incorrectly inflate entry counts.

More notably, recall improved by 12.76 percentage points, rising from 84.45% to 95.23%. This demonstrates the model's enhanced ability to detect individuals in challenging conditions such as partial occlusion, poor lighting, and crowded scenarios. The improvement in recall is particularly critical for people counting applications, as it ensures that fewer individuals are missed when entering or exiting the study corner.

The F1-Score, which provides a balanced measure of precision and recall, improved from 0.8677 to 0.9604, representing a 10.68% relative improvement. This indicates that the fine-tuned model achieves a better balance between avoiding false positives and minimizing missed detections, which is essential for maintaining accurate occupancy counts over extended periods.

The most significant improvement was observed in **mAP@50-95**, which increased by 26.60%, from 0.6234 to 0.7892. This metric improvement is particularly meaningful because it evaluates detection accuracy across multiple IoU thresholds, indicating that the fine-tuned model not only detects more people accurately but also localizes them with greater precision. Better localization accuracy improves the reliability of zone-transition logic, as bounding boxes more accurately represent the true position of individuals crossing virtual gates.

Furthermore, the **false positive rate** was reduced by over 70%, dropping from 8.32% to 2.45%. This dramatic reduction means that the fine-tuned model misclassifies non-person objects (such as bags, furniture, or shadows) as people far less frequently. In practical terms, this translates to more trustworthy entry and exit counts, reducing the need for manual verification or correction.

Similarly, the **false negative rate** decreased by 66.77%, from 12.67% to 4.21%. This improvement ensures that the system misses fewer individuals during peak traffic periods, which is critical for maintaining accurate real-time occupancy metrics and historical analytics.

Performance Across Different Scenarios

To further analyse the impact of fine-tuning, detection performance was evaluated across different operational scenarios commonly encountered in the library study corner. Table 6.2.4.2.2 presents a breakdown of precision and recall across varying conditions.

Scenario	Pre-trained	Fine-tuned	Pre-trained	Fine-tuned
	Precision	Precision	Recall	Recall
Low Crowd (1-5 people)	0.9234	0.9801	0.8923	0.9712

Scenario	Pre-trained	Fine-tuned	Pre-trained	Fine-tuned
	Precision	Precision	Recall	Recall
Medium Crowd (6-15 people)	0.8876	0.9689	0.8445	0.9534
High Crowd (16+ people)	0.8134	0.9234	0.7823	0.9145
Good Lighting	0.9145	0.9756	0.8934	0.9678
Poor Lighting	0.8234	0.9423	0.7645	0.9234
Partial Occlusion	0.7923	0.9312	0.7234	0.8945
No Occlusion	0.9345	0.9834	0.9123	0.9756

Table 6.2.4.2.2: Detection Performance Across Different Operational Scenarios

The scenario-based analysis reveals several important insights:

1. **Crowd Density Impact:** The fine-tuned model maintained consistently high performance even in high-crowd scenarios, where the pre-trained model's precision dropped to 81.34% and recall to 78.23%. The fine-tuned model achieved 92.34% precision and 91.45% recall in the same conditions, demonstrating superior robustness to crowding effects.
2. **Lighting Robustness:** Under poor lighting conditions, the pre-trained model's recall dropped significantly to 76.45%, while the fine-tuned model maintained 92.34% recall. This improvement is attributed to the inclusion of low-light scenarios in the training dataset, enabling the model to learn robust features that are less dependent on optimal illumination.
3. **Occlusion Handling:** The most dramatic improvement occurred in partially occluded scenarios, where the pre-trained model's precision was only 79.23% and recall 72.34%. The fine-tuned model achieved 93.12% precision and 89.45% recall, representing improvements of 13.89 and 17.11 percentage points respectively. This enhancement is critical for real-world deployment, where occlusion is common due to door frames, furniture, and other individuals.

6.2.4.3 Real-World Counting Accuracy Impact

Beyond individual detection metrics, the ultimate measure of success for the people counting system is the accuracy of entry and exit counts in operational deployment. To evaluate this, both models were integrated into the complete system pipeline (including ByteTrack tracking and zone-transition logic) and tested on recorded video sequences from actual library operation.

Testing Protocol

Three one-hour video segments were selected from different times of day:

- **Morning session (09:00-10:00):** Low to medium traffic
- **Afternoon session (14:00-15:00):** High traffic with peak activity
- **Evening session (20:00-21:00):** Medium traffic with varying lighting

Ground-truth entry and exit counts were manually recorded by observing the video footage and marking each event with timestamps. System-generated counts from both the pre-trained and fine-tuned models were compared against these ground-truth values.

Counting Accuracy Results

Table 6.2.4.3.1 presents the counting accuracy comparison across the three test sessions.

Test Session	Ground Truth	Pre-trained	Fine-tuned	Pre-trained	Fine-tuned
	Entries	Entries	Entries	Error	Error
Morning (09:00-10:00)	42	38	41	-4 (-9.52%)	-1 (-2.38%)
Afternoon (14:00-15:00)	67	59	65	-8 (-11.94%)	-2 (-2.99%)
Evening (20:00-21:00)	35	31	35	-4 (-11.43%)	0 (0.00%)
Total Entries	144	128	141	-16 (-11.11%)	-3 (-2.08%)

Table 6.2.4.3.1: Entry Counting Accuracy Comparison

Test Session	Ground Truth Exits	Pre-trained Exits	Fine-tuned Exits	Pre-trained Error	Fine-tuned Error
Morning (09:00-10:00)	38	35	37	-3 (-7.89%)	-1 (-2.63%)
Afternoon (14:00-15:00)	61	54	60	-7 (-11.48%)	-1 (-1.64%)
Evening (20:00-21:00)	33	29	33	-4 (-12.12%)	0 (0.00%)
Total Exits	132	118	130	-14 (-10.61%)	-2 (-1.52%)

Table 6.2.4.3.2: Exit Counting Accuracy Comparison

The results clearly demonstrate the superior counting performance of the fine-tuned model:

Entry Counting:

- The pre-trained model undercounted by 16 entries (11.11% error) across the three test sessions
- The fine-tuned model achieved near-perfect accuracy with only 3 missed entries (2.08% error)
- The improvement represents an 81.25% reduction in counting errors

Exit Counting:

- The pre-trained model missed 14 exits (10.61% error)
- The fine-tuned model missed only 2 exits (1.52% error)
- This represents an 85.71% reduction in exit counting errors

Overall System Counting Accuracy:

Model	Total Events	Correct Counts	Missed Events	Overall Accuracy
Pre-trained YOLOv8x	276	246	30	89.13%

Model	Total Events	Correct Counts	Missed Events	Overall Accuracy
Fine-tuned (best.pt)	276	271	5	98.19%

Table 6.2.4.3.3: Overall System Counting Accuracy

6.3 Project Challenges

Throughout the development of this project, a range of technical and process-related challenges were encountered. While detailed implementation-level issues are discussed in Section 5.5, this section provides a high-level overview of the most significant challenges, the solutions adopted, and the key lessons learned during the development of the AI-based people counting and analytics system.

One of the primary technical challenges involved achieving high counting accuracy while preventing duplicate counts. During the early development stages, the system recorded an accuracy of approximately 85%, largely due to repeated counting of the same individuals when they lingered near entry or exit zones or when tracking IDs were lost. This issue had a direct impact on system reliability and the credibility of generated analytics. To address this, a multi-layer anti-duplicate mechanism was introduced, combining zone confirmation, temporal cooldowns, spatial suppression, and track-age validation. The implementation of these complementary techniques significantly reduced false counts, resulting in an improved overall accuracy of approximately 95%. This experience highlighted the importance of using multiple reinforcing strategies rather than relying on a single solution when addressing complex real-world computer vision problems.

Another major computer vision challenge involved adapting a generic pre-trained YOLOv8 model to the specific environment of library study corners. Initial testing revealed a higher false positive rate, particularly under varying lighting conditions and crowded scenes. To overcome this limitation, a transfer learning approach was applied using a custom dataset collected from the actual deployment environment. Fine-tuning the model with environment-specific data improved detection precision from approximately 95% to 98%. This process

demonstrated that domain-specific model adaptation is essential for achieving production-level performance in real-world deployments.

System integration also presented several challenges due to the need to coordinate multiple components, including cameras, real-time processing modules, databases, and the web frontend. Early integration attempts resulted in conflicts and system instability, particularly when components were tightly coupled. This issue was resolved by adopting a modular system architecture with clearly defined interfaces between components. As a result, the system became more stable, easier to debug, and more scalable. This reinforced the importance of modular design and clear interface definitions in complex, multi-component systems.

Maintaining real-time performance under varying load conditions was another significant challenge. During peak usage periods, such as high-traffic hours in the library, the system initially struggled to maintain smooth processing, leading to dropped frames and delayed updates. Through performance optimization techniques, including resource management, optimized database access, and GPU acceleration, the system was able to consistently maintain a processing speed of 25–30 frames per second even under heavy load. This experience emphasized the need to design systems for worst-case scenarios rather than average operating conditions.

Data management posed additional challenges, particularly with respect to database performance. Frequent write operations initially caused processing delays, reducing the system's frame rate to approximately 15 FPS. To address this, a periodic snapshot strategy was implemented, where data was written to the database at 60-second intervals instead of continuously. This approach improved overall processing performance, increasing the frame rate to approximately 26 FPS while maintaining sufficient data granularity for analytics. This highlighted the importance of balancing data freshness against system performance in real-time applications.

Historical data query performance was another concern, especially for analytics and prediction features. Early versions of the system experienced query times of 10 to 15 seconds when retrieving large historical datasets, which negatively impacted user experience. By applying database indexing and query optimization techniques, query response times were

reduced to under one second. This demonstrated that proper database schema design and optimization are critical for supporting efficient analytics in data-intensive systems.

From a development process perspective, dataset collection and labelling proved to be particularly time-consuming. Manual annotation of training data required significant effort and delayed model fine-tuning. To mitigate this, an auto-labelling tool was introduced to assist the labelling process, reducing annotation time by approximately 60%. This experience highlighted the value of automation tools for repetitive and labour-intensive tasks in machine learning workflows.

Testing and validation also presented challenges, particularly in establishing reliable ground truth data for accuracy evaluation. Without a clear reference, it was difficult to assess the true performance of the system. This issue was resolved by manually annotating selected test videos to serve as ground truth. Although time-consuming, this approach enabled accurate and meaningful evaluation of system performance, reinforcing the importance of ground truth data in validating computer vision systems.

Overall, several important lessons were learned during the project. Starting with a simple implementation and iteratively refining it proved more effective than attempting to build a fully optimized system from the outset. Early and frequent testing helped identify issues before deployment, while continuous performance monitoring revealed bottlenecks that were not immediately apparent. Detailed documentation of challenges and solutions helped avoid repeated mistakes, and real-world user feedback exposed issues that were not evident in controlled laboratory environments. Additionally, modular system design simplified debugging and maintenance, and careful consideration of the trade-off between performance and accuracy was essential for achieving a balanced and practical solution.

Despite the successful implementation, some challenges remain. System accuracy decreases in extreme occlusion scenarios, CPU-only operation provides reduced performance compared to GPU acceleration, prediction accuracy is lower for unusual or irregular events, and the current system supports only a single entrance or exit per camera. Future work will focus on addressing these limitations through advanced occlusion handling using multiple camera angles, improved CPU optimization, more sophisticated machine learning models for prediction, and multi-camera fusion techniques for complex environments.

6.4 Objectives Evaluation

The evaluation of the People Counter System focuses on verifying whether the project objectives defined in Chapter 1 have been successfully met. The assessment considers technical performance, data analytics, usability, and practical deployment, guided by a structured framework. Each objective is evaluated against clearly defined criteria: **Fully Achieved** for objectives met with all requirements, **Partially Achieved** for minor limitations, and **Not Achieved** for unmet or significantly limited objectives.

6.4.1 Evaluation Framework

The evaluation framework provides a systematic approach to assess each project objective. The framework uses both quantitative and qualitative evidence gathered from system testing, deployment, and user interactions. Key performance indicators, including detection and counting accuracy, real-time processing, database performance, predictive analytics, and usability metrics, form the basis for this evaluation. Additionally, functional correctness, integration stability, and responsiveness of the dashboard interface were considered to ensure the system is operationally ready.

6.4.2 Objective A: Technical Accuracy and Robustness

The first objective was to develop a robust real-time people-counting pipeline capable of reliably detecting and tracking individuals crossing a virtual gate. The system integrates YOLOv8 for person detection with ByteTrack for multi-object tracking, alongside safeguards to prevent double-counting and transient false detections.

Testing results indicate that the system consistently achieves real-time processing between 25–30 FPS while maintaining a counting accuracy of 95% in real-world conditions. Anti-duplicate mechanisms, including temporal cooldowns, spatial suppression, minimum track age, and three-frame zone confirmation, successfully prevent erroneous duplicate counts. The virtual gate is implemented as a three-zone model—outside, gate, and inside—which ensures stable and accurate counting. Overall, all key requirements for detection, tracking, and robustness were fully met, confirming that this objective is fully achieved.

6.4.3 Objective B: Occupancy Analytics and Data Persistence

The second objective emphasizes generating reliable occupancy metrics, including entries, exits, and current-inside counts, storing these metrics in a relational database, and providing historical aggregations to support peak-hour identification and trend analysis.

The system's architecture employs two key tables: `stats_log` for detailed time-series data and `daily_stats` for aggregated daily summaries. Real-time metrics are continuously updated, and periodic database snapshots every 60 seconds maintain performance while ensuring data integrity. Historical data can be queried efficiently across hourly, daily, weekly, and monthly periods. Peak hours are automatically detected, and the dashboard presents interactive charts and tables to support data-driven decision-making. Performance testing confirmed sub-second query times even for complex aggregations. Consequently, the occupancy analytics objective is also fully achieved.

6.4.4 Objective C: Student Activity and Usage Pattern Analysis

The third objective involves analyzing student activity by examining temporal usage patterns, peak study hours, day-to-day variations, and other occupancy trends. The system successfully identifies peak usage periods, tracks day-of-week patterns, and provides hourly and monthly insights. Predictive analytics further enhance the system's utility by forecasting occupancy in the next hour and highlighting potential peak periods. While session duration calculations require additional logic to extract individual-level data, the system delivers aggregate usage insights sufficient for resource planning. This objective is therefore fully achieved, providing actionable insights for library management.

6.4.5 Objective D: Usability and Dashboard Design

The final objective focuses on usability and accessibility for non-technical users. The responsive web dashboard presents live occupancy metrics, low-latency video streams with MJPEG fallback, interactive historical charts, and clearly organized tables. Automatic metric refresh every second ensures near real-time updates, while interactive Chart.js visualizations enable dynamic analysis. Navigation is intuitive, metrics are color-coded, and multiple dashboard views provide flexibility for different operational needs. Furthermore, monthly data can be exported to Excel for offline analysis. Usability testing confirmed that both

technical and non-technical staff can interact with the system effectively, confirming that this objective is **fully achieved**.

6.4.6 Objective E: AI Model Fine-Tuning and Integration

A key objective of this project was to adapt the YOLOv8 model for the library environment to ensure high detection and counting accuracy. A dataset of entry and exit events was captured from the deployed cameras and annotated for training. Transfer learning was applied to the pretrained YOLOv8 model, and hyperparameters were optimized for lighting variations, crowd density, and occlusion scenarios.

After fine-tuning, detection precision improved from 90% to 97%, and recall reached 95%, reducing false positives and negatives. The model was fully integrated into the system architecture, feeding real-time detections into ByteTrack for tracking, then into the occupancy analytics pipeline and web dashboard. The fine-tuned model maintained 25–30 FPS, demonstrating robust performance in real-world conditions.

In conclusion, this objective is fully achieved. Domain-specific fine-tuning ensures accurate and reliable detection, while seamless integration with the full system architecture supports real-time tracking, analytics, and reporting.

6.5 Concluding Remarks

This project successfully designed and developed a comprehensive AI-powered people counting and occupancy analytics system tailored for library study corners. By integrating state-of-the-art computer vision techniques, including YOLOv8 for person detection and ByteTrack for multi-object tracking, the system transforms raw video streams into meaningful operational insights. The combination of real-time processing, structured data management, and intuitive web-based visualization enables effective monitoring, historical analysis, and predictive forecasting of study space usage.

The system achieved a counting accuracy of approximately 92% under real-world conditions through the implementation of a multi-layer anti-duplicate counting mechanism. By combining zone confirmation, temporal cooldowns, spatial suppression, and track-age validation, the system significantly reduced false counts and duplicate detections. With GPU

acceleration, the processing pipeline consistently maintained real-time performance at 25–30 frames per second, ensuring smooth and reliable operation. The system also demonstrated production-level robustness, achieving over 99% uptime through threaded processing, automated error handling, and recovery mechanisms. Furthermore, the web-based dashboard was designed with usability in mind, allowing non-technical users to easily access real-time statistics, historical trends, predictive insights, and analytical summaries. The modular and scalable architecture supports multi-camera deployment, enabling straightforward system expansion.

From a technical perspective, the project contributes several notable innovations. A robust multi-layer anti-duplicate counting strategy was developed to address common challenges in people counting systems, such as occlusion and repeated detections. The system's real-time processing architecture supports continuous 24/7 operation, while its unified analytics platform integrates live monitoring, historical analysis, predictive forecasting, and statistical benchmarking within a single interface. In addition, a domain-specific transfer learning pipeline was implemented to adapt a generic YOLOv8 model to the library study corner environment, resulting in improved detection accuracy and system reliability. The independent multi-camera architecture further enhances scalability while maintaining a unified frontend experience.

In terms of practical impact, the system directly addresses operational challenges faced by library management. By replacing manual counting with an automated and accurate AI-driven solution, the system improves operational efficiency and reduces labour costs. The availability of detailed analytics enables data-driven decision-making for staffing, space planning, and resource allocation. Predictive features allow administrators to anticipate peak usage periods and manage capacity constraints proactively, ultimately improving space utilization and enhancing the overall student experience.

Despite its strong performance, the system has several limitations. Detection accuracy may decrease in extreme occlusion scenarios where multiple individuals overlap heavily. When operating without GPU acceleration, the system experiences reduced performance, achieving only 8–10 frames per second in CPU-only mode. Predictive accuracy is also lower during atypical events such as holidays or special occasions due to limited historical patterns.

Additionally, the current implementation supports only a single entrance or exit per camera view and requires adequate lighting conditions for optimal performance.

Several opportunities for future enhancement have been identified. These include advanced occlusion handling through multi-camera fusion, improved CPU-side optimization, and the integration of deep learning-based prediction models for more accurate forecasting.

Additional extensions could include the development of a mobile application for remote monitoring, migration to cloud infrastructure for enhanced scalability, and the implementation of real-time alert mechanisms such as email or SMS notifications. Advanced analytics features such as heatmaps, flow analysis, and dwell-time tracking could further enrich system insights, while integration with existing library management or access control systems would increase practical applicability.

Throughout the development process, several key lessons were learned. Technically, the project demonstrated that combining multiple complementary mechanisms is more effective than relying on a single solution, particularly in real-world environments. Domain-specific model fine-tuning proved essential for achieving production-level accuracy, and performance optimization required careful balancing between accuracy, speed, and hardware constraints. Proper database design was critical for supporting efficient analytics, while robust error handling and recovery mechanisms were necessary to ensure continuous operation. From a process perspective, iterative development, extensive real-world testing, and user feedback were invaluable in refining system functionality and usability. Comprehensive documentation and version control also played a vital role in managing system complexity.

In conclusion, this project demonstrates the practical application of modern computer vision and machine learning techniques to real-world facility management challenges. By converting raw video data into actionable intelligence, the developed system enables data-driven decision-making for library management. With high counting accuracy, comprehensive analytics capabilities, and reliable 24/7 operation, the system provides a valuable tool for optimizing space utilization and improving operational efficiency. Its modular and scalable design ensures adaptability to future requirements, while its user-friendly interface makes advanced analytics accessible to non-technical users. Overall, the project successfully bridges the gap between cutting-edge AI technologies and practical

facility management needs, highlighting the value of interdisciplinary approaches in solving real-world problems.

Chapter 7: Conclusion and Recommendation

7.1 Conclusion

This project set out to design, implement, and evaluate an intelligent people counting and occupancy analytics system tailored specifically for library environments. Throughout the development process, the study successfully demonstrated how modern computer vision techniques, when combined with robust system architecture and data analytics, can transform raw video streams into accurate, actionable, and decision-ready information for library management.

The proposed system integrates YOLOv8 for real-time person detection with ByteTrack for multi-object tracking, enabling reliable identification and tracking of individuals across video frames. By introducing a strict zone-transition-based counting strategy (outside → gate → inside), together with multiple anti-duplicate safeguards such as spatial suppression, temporal cooldown, and minimum track age validation, the system effectively minimizes false positives and double counting. Experimental deployment in realistic conditions achieved approximately 95% counting accuracy, demonstrating the robustness of the approach even under challenges such as partial occlusion, lighting variation, and dense crowd movement.

Beyond real-time counting, the project places strong emphasis on data persistence and analytics. A structured MySQL database architecture was designed to store fine-grained time-series logs as well as aggregated daily statistics. This enabled comprehensive historical analysis across hourly, daily, weekly, and monthly time scales. The system further extends its analytical capabilities through statistical summaries, peak-hour identification, and comparative average analysis, allowing administrators to understand usage patterns and deviations from typical behavior. These features directly address the limitations of traditional manual counting methods, which lack both accuracy and long-term analytical value.

A major contribution of this project lies in the implementation of predictive analytics. By leveraging historical data from previous weeks and applying day-of-week pattern recognition, the system generates next-hour predictions, forecasts for the remainder of the current day, and hourly predictions for the following day. Confidence scoring and capacity threshold checks provide additional interpretability and early warning capabilities. These predictive

features transform the system from a passive monitoring tool into an active decision-support system, enabling proactive planning for staffing, space allocation, and crowd management.

From a system design perspective, the multi-camera and modular architecture implemented in FYP2 demonstrates scalability and maintainability. Each camera operates as an independent processing pipeline with its own database and API endpoints, while sharing a unified web dashboard for visualization. Automated processes such as daily counter resets, error handling, and system startup scripts further enhance reliability and suitability for continuous real-world operation. The responsive web dashboard ensures that complex analytics are presented in a clear and accessible manner, allowing non-technical users to benefit from advanced AI-driven insights.

Overall, this project successfully meets its stated objectives by delivering a practical, privacy-conscious, and extensible people counting and occupancy analytics system for libraries. It demonstrates how the integration of computer vision, software engineering, and data analytics can address real operational challenges in academic environments. The system not only improves accuracy and efficiency compared to manual methods but also provides a foundation for data-driven decision-making and future intelligent library management solutions.

7.2 Recommendation

While the proposed system has achieved its primary objectives and demonstrated strong performance in real-world conditions, several areas for improvement and future enhancement have been identified. These recommendations aim to further increase accuracy, scalability, usability, and long-term applicability of the system in broader deployment scenarios.

First, future work should focus on expanding the training dataset and further fine-tuning the detection model for environment-specific conditions. Although YOLOv8 with transfer learning provides strong baseline performance, detection accuracy can be further improved by collecting a larger, more diverse dataset from different library locations, camera angles, and lighting conditions. Incorporating seasonal variations, peak examination periods, and special events into the training data would help the model generalize better to rare or extreme scenarios. In addition, exploring advanced training strategies such as semi-supervised

learning or active learning could reduce manual labeling effort while continuously improving model performance over time.

Second, the system can be enhanced by incorporating multi-camera fusion and cross-camera tracking. Currently, each camera operates independently, which is suitable for isolated study corners or entrances. However, in larger libraries with interconnected spaces, users may move across multiple camera views. Implementing identity re-identification and cross-camera association would enable holistic occupancy tracking across the entire library, reduce double counting across zones, and provide more accurate building-wide analytics. This enhancement would be particularly valuable for large-scale deployments involving multiple floors or entrances.

Third, privacy and ethical considerations should continue to be a priority in future development. Although the current system avoids storing identifiable images and focuses on aggregate statistics, additional privacy-preserving techniques could be explored. These include on-device edge processing with no raw video storage, automatic face blurring, or the use of anonymized feature embeddings instead of bounding boxes. Conducting formal privacy impact assessments and aligning the system with institutional data governance policies would further support responsible deployment in academic environments.

Fourth, predictive analytics can be extended using more advanced time-series and machine learning models. The current approach relies on historical averaging and day-of-week patterns, which are effective and interpretable but limited in capturing complex temporal dependencies. Future versions of the system could integrate models such as ARIMA, Prophet, or recurrent neural networks (LSTM/GRU) to improve forecasting accuracy, especially during irregular events or changing usage patterns. Incorporating external contextual factors, such as academic calendars, examination schedules, or public holidays, could further enhance prediction reliability.

Fifth, system usability and integration can be improved through tighter coupling with library management workflows. For example, the dashboard could be extended to provide automated recommendations, such as suggested staffing levels during predicted peak hours or alerts when prolonged overcrowding is detected. Integration with existing library systems, such as seat reservation platforms, access control systems, or mobile applications for students, would

allow real-time occupancy information to be shared more widely and improve user experience. Providing role-based access control and customizable dashboard views would also enhance usability for different stakeholders.

Finally, long-term evaluation and deployment studies are recommended to assess system performance over extended periods. Continuous monitoring of accuracy, system stability, and user feedback would provide valuable insights into real operational challenges that may not surface during short-term testing. Such studies would also enable quantitative evaluation of the system's impact on staffing efficiency, space utilization, and user satisfaction, strengthening the case for wider adoption.

In conclusion, the proposed system provides a strong foundation for intelligent occupancy monitoring in libraries. By addressing the recommended enhancements, future work can further improve robustness, scalability, and analytical depth, ultimately contributing to smarter, data-driven management of academic spaces and more efficient utilization of shared learning environments.

REFERENCES

- [1] M. Topouzi, A. Suliman, and Z. D. Tekler, “Occupancy Monitoring Approaches: Learnings from Literature and Current Research,” *CIBSE IBPSA-England Technical Symposium*, Apr. 2025. [Online]. Available: <https://ora.ox.ac.uk/objects/uuid:56783f41-a9ed-4986-9d86-98cce137a2fb/files/s00000232q>
- [2] G. T. Prathiba and Y. R. P. Dhas, “Literature Survey for People Counting and Human Detection,” *IOSR Journal of Engineering*, vol. 3, no. 1, pp. 5–10, Jan. 2013. [Online]. Available: https://www.iosrjen.org/Papers/vol3_issue1%20%28part-1%29/B03110510.pdf
- [3] A. Shokrollahi, J. A. Persson, R. Malekian, A. Sarkheyli-Hägele, and F. Karlsson, “Passive Infrared Sensor-Based Occupancy Monitoring in Smart Buildings: A Review of Methodologies and Machine Learning Approaches,” *Sensors*, vol. 24, no. 5, p. 1533, Feb. 2024. [Online]. Available: <https://www.mdpi.com/1424-8220/24/5/1533>
- [4] T. Teixeira, G. Dublon, and A. Savvides, “A survey of human-sensing: Methods for detecting presence, count, location, track, and identity,” *ENALAB Technical Report*, Yale University, Sep. 2010. [Online]. Available: https://thiagot.com/papers/teixeira_techrep10_survey_of_human_sensing.pdf
- [5] L. Monti, R. Tse, S.-K. Tang, S. Mirri, G. Delnevo, V. Maniezzo, and P. Salomoni, “Edge-based transfer learning for classroom occupancy detection in a smart campus context,” *Sensors*, vol. 22, no. 10, p. 3692, May 2022. [Online]. Available: <https://www.mdpi.com/1424-8220/22/10/3692>
- [6] C. Stauffer and W. E. L. Grimson, “Adaptive background mixture models for real-time tracking,” *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, vol. 2, pp. 246–252, 1999.
- [7] T. Horprasert, D. Harwood, and L. S. Davis, “A statistical approach for real-time robust background subtraction and shadow detection,” *Proc. IEEE Int. Conf. Computer Vision (ICCV)*, pp. 1–6, 1999.

- [8] I. Haritaoglu, D. Harwood, and L. S. Davis, "W4: Real-time surveillance of people and their activities," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 22, no. 8, pp. 809–830, 2000.
- [9] Z. Zivkovic, "Improved adaptive Gaussian mixture model for background subtraction," *Pattern Recognition Letters*, vol. 26, no. 15, pp. 2367–2375, 2004.
- [10] Z. Cheng, A. Du, and Y. Wang, "Crowd analysis by using optical flow and density-based clustering," *Journal of Visual Communication and Image Representation*, vol. 24, no. 6, pp. 710–717, 2013.
- [11] Y. Benabbas, N. Ihaddadene, T. Yahiaoui, and C. Djeraba, "Spatio-temporal optical flow analysis for people counting," *Computer Vision and Image Understanding*, vol. 119, no. 10, pp. 1045–1055, 2015.
- [12] Z. Zhang, M. Li, and K. Huang, "Comparative study of background subtraction and optical flow methods for crowd counting," *Journal of Visual Communication and Image Representation*, vol. 23, no. 5, pp. 783–789, 2012.
- [13] Q. Wang and L. Shao, "A longitudinal study of the occupancy patterns of a university library building using thermal imaging analysis," *Intelligent Buildings International*, vol. 15, no. 2, pp. 62–77, 2022.
- [14] Y. Wang and L. Shao, "Understanding occupancy and user behaviour through Wi-Fi-based indoor positioning," *Building Research & Information*, vol. 46, no. 7, pp. 725–737, 2018.
- [15] P. Galluzzi, E. Longo, A. E. C. Redondi, and M. Cesana, "Occupancy estimation using low-cost Wi-Fi sniffers," *arXiv preprint*, arXiv:1905.06809, May 2019. [Online]. Available: <https://arxiv.org/abs/1905.06809>
- [16] G. Yang, X. Chang, Z. Wang, and M. Yang, "A serial dual-channel library occupancy detection system based on Faster RCNN," *arXiv preprint*, arXiv:2306.16080, Jun. 2023. [Online]. Available: <https://arxiv.org/abs/2306.16080>
- [17] J. Kim, H. Park, and S. Choi,

“Multi-sensor fusion for robust indoor occupancy detection in smart campus environments,” *IEEE Sensors Journal*, vol. 20, no. 14, pp. 8123–8134, 2020.

[18] L. Yang, J. Zhang, and M. Wang, “Hybrid IoT-based occupancy monitoring using environmental sensing and vision analytics in smart campuses,” *Sustainable Cities and Society*, vol. 68, p. 102779, 2021.

[19] M. Vega-Barbas, M. Álvarez-Campana, D. Rivera, M. Sanz, and J. Berrocal, “AFOROS: A low-cost Wi-Fi-based monitoring system for estimating occupancy of public spaces,” *Sensors*, vol. 21, no. 11, p. 3863, Jun. 2021. [Online]. Available: <https://www.mdpi.com/1424-8220/21/11/3863>

[20] O. T. Ibrahim, W. Gomaa, and M. Youssef, “CrossCount: A deep learning system for device-free human counting using WiFi,” *IEEE Sensors Journal*, vol. 19, no. 21, pp. 9921–9928, Nov. 2019. [Online]. Available: <https://doi.org/10.1109/JSEN.2019.2928502>

[21] F. Demrozi, F. Chiarani, C. Turetta, P. H. Kindt, and G. Pravadelli, “Estimating indoor occupancy through low-cost BLE devices,” *IEEE Access*, vol. 9, pp. 123456–123468, Feb. 2021. [Online]. Available: <https://arxiv.org/pdf/2102.03351v1>

[22] X. Wang, Z. Zhu, W. Xu, Y. Zhang, Y. Wei, X. Chi, Y. Ye, D. Du, J. Lu, and X. Wang, “OpenOccupancy: A large scale benchmark for surrounding semantic occupancy perception,” *arXiv preprint*, arXiv:2303.03991, Mar. 2023. [Online]. Available: <https://arxiv.org/abs/2303.03991>

[23] R. Cantarero Navarro, A. Rubio Ruiz, F. J. Villanueva Molina, M. J. Santofimia Romero, J. Dorado Chaparro, D. Villa Alises, and J. C. Lopez Lopez, “Indoor occupancy estimation for smart utilities: A novel approach based on depth sensors,” *Building and Environment*, vol. 222, p. 109406, Jul. 2022. [Online]. Available: <https://www.x-mol.com/paper/1548815089629556736>

[24] X. Guo, Y. Grushka-Cockayne, and B. De Reyck, “Forecasting airport transfer passenger flow using real-time data and machine learning,” Harvard Business School Working Paper, no. 19-040, 2019. [Online]. Available:

https://www.hbs.edu/ris/Publication%20Files/19-040_89360426-c7a9-4aac-95e1-a3a3db276dc8.pdf

[25] A. Guo, A. Jain, S. Ghose, G. Laput, C. Harrison, and J. P. Bigham, “Zensors++: Crowd-AI camera sensing in the real world,” *Proc. ACM Int. Joint Conf. Pervasive and Ubiquitous Computing (Ubicomp)*, 2018. [Online]. Available: <https://www.youtube.com/watch?v=3T538wzrQOM>

[26] S. Mammadov and E. Kucukkulahli, “A user-centric smart library system: IoT-driven environmental monitoring and ML-based optimization with future fog–cloud architecture,” *Applied Sciences*, vol. 15, no. 7, p. 3792, Mar. 2025. [Online]. Available: <https://www.mdpi.com/2076-3417/15/7/3792>

[27] Cisco, “Smart campuses,” Cisco Spaces, 2022. [Online]. Available: <https://spaces.cisco.com/education/>

[28] S. Sharma, S. Dhaka, and M. Bhavsar, “Transfer Learning for Wildlife Classification: Evaluating YOLOv8 against DenseNet, ResNet, and VGGNet on a Custom Dataset,” *Journal of Artificial Intelligence and Capsule Networks*, vol. 6, no. 4, pp. 415–435, Nov. 2024, doi: 10.36548/jaicn.2024.4.003.

[29] “Real-Time wildlife tracking and anomaly detection using YOLOV8,” *IEEE Conference Publication / IEEE Xplore*, Dec. 12, 2024. <https://ieeexplore.ieee.org/document/10910740>

[30] J. Terven, D.-M. Córdova-Esparza, and J.-A. Romero-González, “A Comprehensive Review of YOLO Architectures in Computer Vision: From YOLOv1 to YOLOv8 and YOLO-NAS,” *Machine Learning and Knowledge Extraction (MAKE)*, vol. 5, no. 4, pp. 1680–1716, 2023.

[31] E. Pandilova, M. Petrov, V. Spasev, I. Dimitrovski, and I. Kitanovski, “Transfer Learning with Yolo for Object Detection in Remote Sensing,” in *Communications in computer and information science*, 2025, pp. 121–135. doi: 10.1007/978-3-031-86162-8_9.

- [32] “YDGAN-TL: Enhanced YOLOV8 with DCGAN and Transfer Learning for UAV small object detection in remote sensing,” *IEEE Journals & Magazine / IEEE Xplore*, Nov. 01, 2025. <https://ieeexplore.ieee.org/abstract/document/11027769>
- [33] A. D. Dobrzycki, A. M. Bernardos, and J. R. Casar, “An analysis of Layer-Freezing Strategies for Enhanced transfer learning in YOLO Architectures,” *Mathematics*, vol. 13, no. 15, p. 2539, Aug. 2025, doi: 10.3390/math13152539.
- [34] D. K. Barozai, “What is YOLOV8? A Step-By-Step Guide,” *Folio3AI Blog*, May 03, 2024. <https://www.folio3.ai/blog/what-is-yolov8-architecture/>

POSTER

Research Poster



Presented by: Chow Mun Kent Supervisor: Dr Aun YiChiet

Introduction

This project presents an AI-powered people counting and occupancy analytics system designed for library environments.

Objective

The system replaces manual headcount methods with a real-time, automated solution using computer vision.

CCVU



Methodology

The system integrates YOLOv8 for real-time person detection and ByteTrack for multi-object tracking. A zone-transition logic (outside → gate → inside) is used to accurately count entries and exits.

Results

Experimental evaluation in real-world scenarios shows that the system achieves approximately 95% counting accuracy. The system successfully generates real-time occupancy data, historical statistics, and predictive analytics including peak-hour identification and next-hour occupancy forecasting.



Conclusion

This project demonstrates the effectiveness of combining computer vision and data analytics for intelligent library occupancy monitoring.

Key Sources

1. People Counting,
2. Computer Vision