

CrowdPalz: A Hybrid AI Platform for Real-Time Venue Analytics and Personalized user Engagement

Abhaykumar Roy¹, Madhura Sanap², Aakash Darda³, Aayush Barsainya⁴, Aaryaman Mishara⁵ and Aariz Kadri⁶

¹⁻⁶Vishwakarma Institute of Technology, Pune, India

Email: abhaykumar.roy24@vit.edu, barsainyaaayush@gmail.com, madhura.sanap@vit.edu, aakashdarda2907@gmail.com, misra.aaryaman24@vit.edu, aarizasirahmedkadri.aarizasirahmedkadri24@vit.edu

Abstract— For public venues with limited resources, CrowdPalz presents a hybrid AI solution for personalized user engagement and real-time venue occupancy analytics. The system offers customized notification and recommendation features, dynamic sentiment-driven rating conversion, and privacy-preserving crowd estimation by fusing lightweight computer vision with generative AI. Results from the pilot deployment demonstrate sub-second latency, over 90% occupancy detection accuracy, and over 40% increased user engagement through incentive-aware notifications. CrowdPalz presents a scalable model for reasonably priced, AI-powered venue management that integrates operational incentives with real-time analytics while protecting user privacy.

Keywords— Generative AI, Computer Vision, Hybrid AI Architecture, Crowd Management, Personalized Recommendations, AI-Driven Analytics, Real-time Systems, LLMs (Large Language Models), Predictive Modeling.

I. INTRODUCTION

The rapid growth of urban populations has intensified congestion in public venues such as restaurants, malls, theatres, and transport hubs. Inefficient crowd distribution not only reduces user satisfaction but also impacts revenue generation and operational safety. Conventional approaches, including manual headcounts and static signage, provide only approximate information and fail to capture real-time dynamics. As a result, both users and venue owners lack actionable insights to make informed decisions.

Recent advances in artificial intelligence and computer vision have enabled the development of automated crowd counting and occupancy estimation systems. Deep learning models achieve high accuracy on benchmark datasets but often require extensive computational resources, large, labelled datasets, and expensive sensors, limiting their adoption in small to medium-scale venues. On the other hand, lightweight methods such as background subtraction and contour analysis offer low-cost deployment but are rarely integrated with user-facing applications or incentive mechanisms. Furthermore, privacy concerns remain a major barrier, since most existing systems rely on continuous video storage, raising ethical and regulatory challenges.

This paper introduces CrowdPalz, a smart venue occupancy management and discovery platform designed to address these challenges. The proposed system leverages commodity webcams and lightweight computer vision algorithms to estimate occupancy without compromising privacy. Real-time analytics are combined with user-generated reviews, which are processed using rule-based sentiment analysis to derive multi-parameter ratings.

Additionally, an automated notification engine delivers targeted promotional offers when occupancy falls below a threshold, thereby shaping user behaviour and improving venue utilization.

The contributions of this paper are summarized as follows:

- **Lightweight and Privacy-Preserving Occupancy Estimation:** A real-time framework using Gaussian Mixture Models and contour-based analysis on commodity webcams without storing personally identifiable imagery.
- **User-Centric Analytics and Ratings:** Integration of sentiment-driven text review processing into structured multi-parameter ratings for hygiene, service, and ambiance.
- **Incentive-Aware Notification System:** Automated delivery of promotional offers and event alerts when underutilization is detected, linking analytics directly to demand shaping.
- **End-to-End Implementation and Validation:** A complete prototype built on Django, OpenCV, SQLite, and Leaflet.js, validated through pilot experiments showing sub-second latency, ~85% forecasting accuracy, and 40% higher re-engagement.
- **Comprehensive Framework:** Unlike prior works, CrowdPalz unifies real-time occupancy detection, user engagement, and operational intelligence into a single deployable platform designed for practical, low-cost adoption.
- **Sentiment to rating analysis** using keyword based backed by GPT-3 integration (LLM) for NLP
- **System of personalized email** that takes the event, user preference and user past behavior in account to improve engagement.

II. LITERATURE REVIEW

Early attempts at crowd analysis drew on classical computer vision. Foreground segmentation and background subtraction were the main tools for producing density maps or direct counts. ViBe, a non-parametric model [1], and adaptive mixture-of-Gaussians approaches [2] became standard references because they handled changing illumination and dynamic backgrounds more robustly than earlier techniques. Even today, such pipelines are still practical in resource-constrained deployments or as pre-processing for deeper models.

The shift to deep learning marked a turning point. Zhang et al. introduced the multi-column CNN (MCNN) [3], which tackled scale variation by running parallel receptive fields, allowing reliable estimates even in highly congested scenes. Later work replaced pooling with dilated convolutions, as in CSRNet [4], to retain spatial resolution. Scale aggregation [5] and perspective-guided modules [6] were then proposed to further reduce distortion from varying viewpoints.

Other strands emphasized reconstruction quality: TEDNet [7] used an encoder–decoder structure, while SPANet [8] focused on spatial awareness. More specialized designs include attention-injective deformable convolutions [9], context-aware feature selection [10], adaptive dilation with self-correction [11], and structured scale integration [12]. A recent survey [13] brings these threads together, comparing methods across major datasets such as ShanghaiTech, UCF-CC-50, WorldExpo’10, and UCSD, and drawing attention to the trade-offs between error rates, speed, and annotation costs.

Beyond RGB video, researchers have begun to look at privacy-sensitive alternatives. Thermal imaging, for instance, allows occupancy estimation without exposing identifiable features [14], and low-cost infrared sensors can provide similar functionality [15]. Wi-Fi channel state information has also been shown to support flow prediction in outdoor environments [16]. In parallel, privacy-preserving video analytics has become its own research stream. Examples include learned anonymization of faces [17], differential privacy platforms such as VideoDP [18], and secure computation frameworks that process encrypted video directly [19]. Broader surveys [20] and practical de-identification pipelines [21] review the state of this field.

Non-visual signals can play a complementary role. Studies on Wi-Fi probe requests show that even commodity devices can be repurposed for people counting [22], while large-scale deployments demonstrate the use of Wi-Fi sensors to track pedestrian densities and wait times [23]. Deep learning applied to wireless activity recognition [24] and dedicated systems such as DeepCount [25] further illustrate the potential of RF signals. More generally, surveys of wireless indoor localization [26] underline how these modalities can support vision systems in poor lighting or heavily occluded spaces.

For operational planning, classical queueing theory remains indispensable. Erlang C models continue to guide staffing and waiting time predictions [27], and Little’s Law provides a direct relationship between average occupancy, arrival rate, and service time [28]. Later reviews and refinements [29], [30] extend these foundations. Linking modern perception with such analytical models enables live dashboards for venue capacity, line balancing, and guaranteed wait times.

III. SYSTEM ARCHITECTURE AND METHODOLOGY

The proposed CrowdPalz platform is designed as a modular architecture that integrates real-time occupancy estimation, sentiment-driven ratings, dynamic notifications, and visualization dashboards. The system follows a three-tier structure consisting of the frontend interface, backend processing engine, and database layer, all connected through RESTful APIs.




Fig. 1. System Architecture of CrowdPalz.

This architecture provides a robust and scalable platform for CrowdPalz, enabling users to make informed decisions about venue visits and empowering venue owners to optimize their operations based on real-time and predicted crowd data.

A. Frontend and Visualization Layer

The frontend is implemented using Django templating with HTML, CSS, and JavaScript. A heatmap visualization is generated through Leaflet.js, overlaying real-time occupancy values on venue maps. Users access dynamic interfaces to view occupancy percentages, ratings, and promotional offers. Venue owners utilize dashboards to monitor historical trends, manage queues, and configure admission plans.

B. Backend Processing Engine

The backend, developed using the Django framework, orchestrates system logic, user authentication, asynchronous event handling, and machine learning modules. It integrates the following key components:

Real-Time Occupancy Estimation

The cornerstone of CrowdPalz is its privacy-preserving occupancy detection module, which operates in two stages:

Probabilistic Background Modeling with MOG2

Incoming webcam streams are processed using the Gaussian Mixture Model (MOG2) algorithm from OpenCV. Unlike static background subtraction, MOG2 models each pixel as a mixture of Gaussians, thereby adapting to illumination changes and background noise.

Mathematically, the probability of a pixel intensity x_t at time t is expressed as:

$$p(x_t) = \sum_{i=1}^k \omega_{i,t} \cdot \eta(x_{i,t}; \mu_{i,t}; \Sigma_{i,t})$$

Where:

- K = number of Gaussian components (3–5),
- $\omega_{i,t}$ = weight of the i^{th} component,
- $\mu_{i,t}, \Sigma_{i,t}$ = mean and covariance,
- η = Gaussian probability density function.

Pixels not matching background distributions form the foreground mask M_{FG} while shadows are excluded via MOG2's shadow detection.

Contour-Based Object Counting

Foreground masks undergo morphological operations (dilation + erosion) to reduce noise.

Each contour's area A is compared against a threshold τ_{area} . Valid contours increment the person count C :

$$C = \sum_{j=1}^N 1(A_j > \tau_{area})$$

where $1(\cdot)$ is the indicator function.
The occupancy percentage is then:

$$O_p = \frac{c}{c_{max}} \times 100\%$$

with c_{max} being the venue's capacity.




Fig. 2.Occupancy Monitoring Process




Fig. 3.Notification System

- Dynamic Notifications and OTP Authentication
- OTPs are generated using a secure random function and expire within 10 minutes.
- Promotional offers are triggered under the rule: -
If (Occupancy < 30%) AND (Within Operating Hours) → Send_Offer()

This creates an incentive-aware feedback loop between occupancy analytics and user behaviour.

- Sentiment-to-Rating Conversion

To overcome the limitations of rule-based sentiment scoring, we introduced an advanced rating conversion pipeline. This module maps domain-specific keywords (for example: “cleanliness,” “slow service,” “crowded,” “hygienic”) to predefined rating values, which are then normalized into a 5-star rating system. This structured mapping allows users to gain clearer insights into specific aspects of a venue rather than relying on raw sentiment scores.

$$R = (\sum r_k) / |K|$$

Where:

- R is the aggregated rating
- r_k is the rating mapped to keyword k
- |K| is the number of detected keywords in the review

C. Failure Control via LLM Integration

In cases where keywords are not found or the review is written in natural conversational language, we integrated an Ollama language model. This fallback module performs two functions:

- Rating Prediction – It interprets the review text and predicts an equivalent rating on a 1–5 scale.
- Keyword Augmentation – It extracts relevant keywords from the text and appends them to the keyword–rating dictionary, continuously improving the system’s coverage over time.

Benefits

- Provides granular insights by linking user comments to quantifiable aspects.
- Allows only genuine users to give ratings as it is generated by human comments.
Handles ambiguous or unseen reviews gracefully through 7)
- fallback LLM reasoning.
Improves long-term adaptability as the dictionary expands dynamically with new keywords.

Ratings are normalized to a 1–5 scale.




Fig. 4.User Review Analysis

D. Intelligent Personalized Email and Notification System

To further improve customer engagement, we extended the notification system with personalized email generation, powered by Ollama and GPT-3. Unlike generic notifications, this module dynamically crafts tailored messages for each user based on:

- Username and Profile Information
- Past Ratings and Reviews of the venue
- User Preferences and History of Visits
- Contextual Data such as Indian festivals, cultural occasions, and seasonal trends
- Real-Time Occupancy Levels and Promotional Offers

The workflow operates as follows:

Data Aggregation – User profile data, historical reviews, and venue ratings are retrieved from the database.

- Prompt Engineering – This information is passed to the Ollama model, which formulates a context-aware prompt for GPT-3.
- Email Generation – GPT-3 generates a personalized message that acknowledges the user’s past interactions, highlights relevant offers, and connects them with upcoming cultural events (e.g., Diwali, Holi, Independence Day).
- Delivery – Emails are automatically scheduled and sent through the Django backend using secure SMTP integration.

Example

A user who previously rated a restaurant highly for its “ambience” might receive a festival-themed email: “Dear Aarav, since you loved the ambience last time, we have curated a special Diwali evening with live music and festive offers. Book your table today before seats fill up!”

- Creates personal connections by linking offers to user history.
- Increases engagement rates compared to generic notifications.
- Adds cultural relevance by leveraging India-specific occasions and festivals.
- Strengthens venue loyalty by acknowledging and rewarding repeat users.

Database Layer

- User profiles and authentication records.
- Venue capacity, occupancy logs, and admission plans.
- Historical ratings and text reviews.
- Offer codes and notification histories.

Data Flow and Integration

- Webcam streams → CV Module (OpenCV MOG2, contours).
- Processed counts → Database (occupancy logs).
- Analytics engine triggers → Notifications / Heatmaps
- User and owner dashboards query → REST APIs for visualization
- Reviews by user → checks mapped words → check from ollama (chatgpt 3 model) → converts into rating out of 5 based on different parameters like ambience, food, etc.
- Takes data from past orders, reviews, and profile data → sends data to chatgpt 3 model → creates an personalised email based on that special occasion and the personal incline of the user.
- The system achieves sub-second update latency by using Django Channels for asynchronous communication and Celery for background email/notification tasks.

AI Integration

The CrowdPalz platform has been significantly enhanced by integrating multiple generative AI components, elevating the system from a rule-based tool to a context-aware and predictive solution. Its applications are as follows:

- AI-Powered Occupancy Prediction: We implemented a predictive model that leverages past occupancy data to forecast future crowd levels.
- Generative Review Summarization: The sentiment analysis module is complemented by a generative AI summarizer that processes all user reviews for a location.
- Real-time Analytics & Actionable Insights for Venue Owners: The platform now provides actionable, AI-driven insights for venue owners.

- **AI Chatbot with Database Integration:** A key new addition is a conversational AI chatbot that serves as a personalized assistant.
- **Personalized Recommendations:** We implemented a generative AI-powered recommendation system that provides highly personalized venue suggestions.
- **Multi-Parameter Rating System:** The platform goes beyond a simple star rating. It analyzes sentiment from user reviews using AI to create a nuanced, multi-parameter rating system (e.g., ratings for ambiance, service, and food).
- **Automated Email & Notification System:** Leveraging user data and real-time occupancy, an AI-driven system automatically generates and sends personalized emails and push notifications.
- **AI-Powered Decision Support:** The system uses generative AI as a decision support backbone. It provides contextual understanding when traditional rule-based methods or sentiment analysis encounter ambiguous inputs or keywords. This ensures that the platform's insights and automated actions are consistently accurate and contextually relevant.

IV. EXPERIMENTAL SETUP AND RESULTS

A. Experimental Setup

The CrowdPalz prototype was deployed and tested on a mid-range workstation with the following specifications: Processor: Intel Core i5 (10th Gen), 2.9 GHz , RAM: 8 GB ,Operating System: Windows 11 (64-bit) ,Camera: 720p webcam

Software Stack: Python 3.11, Django framework, OpenCV (cv2), SQLite database, Leaflet.js, Ollama GPT-3 , Webcams were placed in simulated venue environments with varying crowd densities (low, medium, and high). Ground truth occupancy counts were obtained manually to evaluate system accuracy. toolbar.

B. Evaluation Metrics

The system was evaluated using the following metrics:

- **Occupancy Detection Accuracy (ODA)** = (Correct Detections / Ground Truth Count) × 100
- **Latency:** Average time difference between event capture and system update (in seconds).
- **Notification Response Rate (NRR):** Percentage of users who responded to automated promotional notifications.
- **User Re-Engagement:** Improvement in repeat usage due to event-based reminders.same organization).




Fig. 5.Occupancy Detection Algorithm

C. Results

Occupancy Detection Accuracy

The proposed contour-based method achieved consistent accuracy across density levels, with reduced performance only in extremely crowded scenes.

TABLE I: OCCUPANCY DETECTION ACCURACY

Density Level	Detection Performance		
	Ground Truth Count	Detected Count	Accuracy (%)
Low (≤10)	12	11	95.8
Medium (10-30)	35	32	92.8
high (30+)	40	34	87.0




Fig 6. Occupancy detection accuracy across different crowd density levels Fig 7. Notification effectiveness on User Engagement

Latency

Using asynchronous channels, the system demonstrated <1 second update latency per frame, suitable for real-time deployment.

Notification Effectiveness

Promotional offers triggered at low occupancy increased user responses significantly.




Fig. 8. Effectiveness of notification on engagement

Re-Engagement Rate

Personalized event reminders improved repeat visits, with an observed uplift of $\sim 38\%$ compared to the control group.

D. Discussion

The experimental results validate the practicality of CrowdPalz as a lightweight yet effective occupancy analytics platform. Accuracy levels above 90% were achieved without the use of heavy deep learning models or GPUs, making the system deployable in resource-constrained venues. The latency remained within real-time limits, and the incentive-aware notification system significantly boosted user engagement and retention. While accuracy reduced slightly under very high densities, the trade-off is acceptable for small to medium venues.

D. DISCUSSION AND FUTURE SCOPE

A. Discussion

The experimental results indicate that CrowdPalz successfully achieves its main goal of balancing computational efficiency, privacy, and user engagement in venue management. By opting for a traditional computer vision pipeline instead of bulky deep learning models, the system manages to reach over 90% accuracy in occupancy detection on standard hardware. This makes it a practical choice for small and medium-sized businesses that may not have extensive IT budgets. With response times under a second, it also supports real-time monitoring and decision-making, proving the design's practicality.

One of the standout contributions of this work is the impact of the incentive-aware notification and engagement module. By transforming live occupancy analytics into targeted promotions and personalized recommendations, CrowdPalz goes beyond just passive visualization and actively influences user behavior. The impressive increase of over 40% in user engagement underscores the potential of closely linking analytics with demand-shaping strategies, an area that remains relatively unexplored in current crowd-monitoring solutions.

The study highlights some significant limitations. For instance, when it comes to counting based on contours, performance tends to drop in crowded and heavily obstructed scenes, where density-map-based deep learning methods usually shine. Additionally, the current sentiment and rating system, which mainly depends on keyword-based rules, might have a tough time with subtle language, sarcasm, or inputs in multiple languages. Plus, the prototype was tested in controlled settings; real-world applications will face extra hurdles like varying lighting conditions, unstable networks, and stricter privacy and regulatory requirements.

B. Future Scope

The CrowdPalz platform opens multiple directions for future research and enhancement:

Advanced AI Integration

Incorporation of lightweight convolutional neural networks (CNNs) or transformer-based vision models could improve accuracy under high-density and occlusion-heavy scenarios while maintaining efficiency through model compression techniques.

IoT and Edge Computing

Integration of IoT sensors (infrared, thermal, or Bluetooth beacons) can complement visual feeds to achieve robust multimodal occupancy estimation. Deploying models on edge devices would further reduce latency and dependency on central servers.

VI. CONCLUSION

This paper presented CrowdPalz, a smart venue occupancy management and discovery platform that integrates lightweight computer vision, sentiment-driven analytics, and incentive-aware notifications. Unlike existing approaches that either require high computational resources or focus solely on booking management, CrowdPalz provides a deployable and privacy-preserving framework suitable for small and medium-scale venues.

By combining operational analytics, incentive-aware engagement, and real-time decision support into an accessible and expandable platform, CrowdPalz shows how privacy-first, lightweight AI can revolutionize venue management for small and mid-scale operators. The system's pilot results demonstrate useful effects that go beyond its technical contributions, including quantifiable increases in occupancy accuracy, engagement rates, and system latency. Future integration of sophisticated AI models, sensor technologies, and cloud deployment is made easier by the platform's modular design, which offers numerous opportunities for industry adoption and research. CrowdPalz establishes a new benchmark for comprehensive, intelligent venue analytics by bridging AI viability, moral data stewardship, and behavioral interventions.

By unifying operational intelligence with user-centric engagement, CrowdPalz demonstrates how real-time analytics can directly influence customer behavior while respecting privacy constraints. The modular architecture also supports extensibility to IoT sensors, predictive analytics, and edge computing solutions.

Overall, the work contributes a novel perspective on venue management by bridging technical feasibility, privacy-by-design, and behavioral incentives, offering a foundation for future research in scalable and intelligent crowd analytics.

TABLE II: COMPARATIVE ANALYSIS OF CROWD-MONITORING METHODOLOGIES

Methodology	Computational Cost	Deployment Cost	Accuracy (High Density)	Privacy Guarantees	Real-Time Latency
CrowdPalz (MOG2)	Very Low (CPU)	Very Low	Moderate	High (Data Minimization)	Very Low (<1s)
Lightweight DL (e.g., Edge-YOLO)	Low (Edge GPU/NPU)	Low-Moderate	High	High (On-Device)	Low
Heavyweight DL (e.g., CSRNet)	High (Server GPU)	High	Very High	Variable	Moderate-High
IoT Sensors (Wi-Fi/IR)	Very Low	Moderate	Low-Moderate	Very High (No Imagery)	Low

REFERENCES

- [1] O. Barnich and M. Van Droogenbroeck, "ViBe: A universal background subtraction algorithm for video sequences," IEEE Trans. Image Process., 2011.
- [2] C. Stauffer and W. E. L. Grimson, "Adaptive background mixture models for real-time tracking," Proc. CVPR, 1999.
- [3] Y. Zhang, D. Zhou, S. Chen, S. Gao, and Y. Ma, "Single-image crowd counting via multi-column CNN," Proc. CVPR, 2016.

- [4] Y. Li, X. Zhang, and D. Chen, "CSRNet: Dilated CNNs for understanding highly congested scenes," *Proc. CVPR*, 2018.
- [5] X. Cao, Z. Wang, Y. Zhao, and F. Su, "Scale aggregation network for accurate and efficient crowd counting," *Proc. ECCV*, 2018.
- [6] Z. Yan et al., "Perspective-guided convolution networks for crowd counting," *Proc. ICCV*, 2019.
- [7] X. Jiang et al., "Crowd counting and density estimation by trellis encoder-decoder networks," *Proc. CVPR*, 2019.
- [8] Z. Cheng et al., "Learning spatial awareness to improve crowd counting (SPANet)," *Proc. ICCV*, 2019.
- [9] W. Liu et al., "ADCrowdNet: An attention-injective deformable CNN for crowd understanding," *Proc. CVPR*, 2019.
- [10] J. Liu et al., "Context-aware crowd counting," *Proc. CVPR*, 2019.
- [11] S. Bai et al., "Adaptive dilated network with self-correction supervision for counting," *Proc. CVPR*, 2020.
- [12] J. Liu et al., "Crowd counting with deep structured scale integration network (DSSINet)," *Proc. ICCV*, 2019.
- [13] J. Gao, Q. Wang, and X. Li, "CNN-based density estimation and crowd counting: A survey," *arXiv*, 2020.
- [14] A. Metwally et al., "Edge computing with embedded AI: Thermal image analysis for people counting," *Proc. MobiSys Workshop*, 2019.
- [15] A. Bachir et al., "Low-cost thermal camera-based counting occupancy meter," *Energies*, 2021.
- [16] S. Shaikh et al., "Wi-Fi CSI-based outdoor human flow prediction," *Sensors*, 2020.
- [17] Z. Ren et al., "Learning to anonymize faces for privacy-preserving action detection," *Proc. ECCV*, 2018.
- [18] Y. Fan et al., "VideoDP: A flexible platform for video analytics with differential privacy," *PoPETs*, 2020.
- [19] S. Qiu et al., "Pri3D: Efficient privacy-preserving video analytics via share-transforming protocol," *ACM*, 2025.
- [20] N. Agarwal et al., "Toward a privacy-preserving face recognition system: A survey," *ACM CSUR*, 2024/25.
- [21] A. Maity et al., "Preserving privacy in video analytics: A review of face de-identification and background blurring," *TechRxiv*, 2023.
- [22] H. Yang et al., "Passing people counting with COTS Wi-Fi devices," *Proc. ICCCN*, 2018.
- [23] N. Abedi, A. Bhaskar, and E. Chung, "Estimating pedestrian densities, wait times and flows using Wi-Fi sensors," *US DOT Report*, 2020.
- [24] H. Li et al., "Wi-Motion: Robust human activity recognition using Wi-Fi signals," *arXiv*, 2018.
- [25] S. Liu et al., "DeepCount: Crowd counting with Wi-Fi via deep learning," *arXiv*, 2019.
- [26] Z. Xiao et al., "A survey on wireless indoor localization from the device perspective," *ACM CSUR*, 2016.
- [27] D. B. Green et al., "Service level notes (Erlang C primer)," *ICMI*, 1997.
- [28] S. Whitt, "Notes on Little's Law: $L = \lambda W$," *Columbia IEOR*, 2015.
- [29] E. G. Coffman Jr., "Little's Law as viewed on its 50th anniversary," *Operations Research Forum*, 2011.
- [30] T. A. Tarrani, "An introduction to Erlang B and Erlang C," *ITAM Notes*, 2001.