International Journal on Recent and Innovation Trends in Computing and Communication
ISSN: 2321-8169 Volume: 10 Issue: 11
***Article Received:** 25 July 2022 **Revised:** 12 September 2022 **Accepted:** 30 November 2022*

An AI-Driven Spatiotemporal Crowd Orchestration Platform for Large-Scale Theme Parks: Hybrid Machine Learning, Behavioural Modelling, and Real-Time Decisioning for Safe and Efficient Guest Flow

Suresh Chaganti
Architect Data & ML OPS

ABSTRACT: This paper describes an AI-based spatiotemporal crowd coordination system that can be used in large theme parks. The system relies on machine learning, behavior modeling, and real-time decision rules to forecast congestion and direct the flow of guests as well as keep the density safe. It processes real-time signals like entry surges, attraction queues and movement of pathways to predict the formation of the hotspots. The outcomes of the actual implementations indicate a 32-47 percent congestion decrease, 28 percent queue stability, and adherence to social-distancing thresholds when at peak times. The system enhances safety, flow efficiency and general guest experience providing a model that can be used elsewhere in high-density public areas.

KEYWORDS: Guest Flow, Spatiotemporal, Hybrid, Orchestration, Machine Learning, AI, Decision, Behavioural Modelling

I. INTRODUCTION

Big theme parks have complicated patterns of people movement that are determined by the attraction cycles, weather, events, and behavior of visitors. These obstacles were worsened during the COVID-19 times when there were strict distancing and density restrictions. The conventional crowd-control techniques are too slow and in many cases personnel are subject to manual judgement. The current paper proposes a crowd orchestration platform, based on AI and capable of predicting crowd formation, proposing early intervention, and providing real-time guidance to guests. The system relies on forecasting models, behavioural patterns and simulation engine to enhance efficiency in flow and safety. It sets out to demonstrate how crowd management based on data may facilitate the efficient functioning and enhanced customer experience in very large crowds.

II. RELATED WORKS

Machine Learning for Modern Crowd Management

Recent studies indicate that machine learning has taken the center stage in enhancing massive management of crowds particularly in high-density areas like cities, transport centers, and major events. According to a detailed systematic literature review of the literature published during 2010-21, it is found that either supervised or unsupervised deep learning techniques are common to learn crowd flow, identify anomalies, and forecast future threats with high levels of accuracy [1].

Such methods are usually used on the scenes photographed by cameras, but one of the significant missing links is the use of social-media data that can be used to get more valuable information regarding movements, intentions, and feelings in high volumes. This disjunction also encourages more recent systems of AI to incorporate multimedia data streams particularly in occurrences where human participants act in an unpredictable manner.

The issue of smart crowd control gained more acuity during the COVID 19 pandemic. Numerous mass gatherings did not adhere to the measures of distancing, as such locations became threat-prone areas of spreading the virus [2]. In this regard, scholars suggested machine-learning-driven models consisting of an event-scheduling task, behavioral forecasting, and real-time safety verifications.

Some of the restrictions implemented by the proposed algorithms include mask compliance and safe distancing.

They also incorporate validation measures in order to denote the effectiveness of overtime crowd control strategies [2]. Such attempts prove the usefulness of the AI-based orchestration systems in the case of crisis-driven processes, when the risk level may fluctuate quickly and where the conventional processes of manual monitoring are not adequate.

Machine learning too has been represented in the situation of city scale events where there is a deviant crowd behavior that is not part of the usual daily routine. As an illustration, DeepUrbanEvent is the online deep learning network which forecasts an hour ago the movement of crowds through the examination of the past hour density patterns of the city [3].

It applies recurrent neural network models to sequential data such as distribution of traffic and crowds. These systems have been experimented to be better than traditional models in dynamic city situations. This type of machine learning has been useful to the requirement of spatiotemporal and real time forecasting in places like theme parks where the behavior of guests constantly changes as per the schedule of attractions, weather and unforeseen congestion.

Deep learning is observed to aid in long-term forecasting of the crowd flow in complicated settings. In one work, CAGE, a representation system that is integrated with SegNet based systems became known; it is capable of forecasting the movement of crowds in large regions at any time instantaneously [5].

This system is efficient to manage many of the simulation results in real time, which is necessary in an environment that requires an operator to test a variety of what-if scenarios before a decision-making process. The given research directions give good evidence about the use of high-order forecasting models in orchestrating crowds.

Spatiotemporal Modeling Techniques

Proactive crowd management requires proper forecasting. Crowd prediction innovations have gone beyond short time estimation to the multi-days or even weekly predictions. By way of example, a recent study presented a way of forecasting levels of crowds more than a week ahead of large occasions by supplying mobility logs, transit search information, and situational characteristics, including day-of-week patterns [6].

Their multi-task Poisson regression model is geographically complemented and it gains up to 42% as compared to previous models. The development is significant to the mega theme parks, which rely on the vacation times, ticket reservation, weather conditions, and advertisements to determine the number of people arriving. Early prediction allows the operators time to move personnel, reassign attractions and reformulate routing maps.

Besides highly sophisticated statistical models, digital twin technologies become important constituents of massive event planning. The 3D Digital Twin platform, combined with real-time data, and mobile phone signals was used in the project Scheveningen Crowd Safety Manager. The system assists in operations as well as planning with the help of risk estimation, prediction of traffic-flow and forecast of multi-day crowds [4].

The use of machine learning algorithms like XGBoost was tested and attained very good predictions in the complicated coastal environments. It testifies to the increasing importance of integrating virtual simulation systems and real-life sources of intelligence. Digital twins could be utilized in a theme-park setting to test new routing policies, entertainment schedules, and attraction capacities and roll them out.

Risk-oriented forecasting is another area that has been extended to spatiotemporal modeling. To illustrate, there is machine learning and georeferenced biometric information of wearables to identify a panic event in real time [7]. Through the stress level prediction using Gaussian SVM classifiers and mapping of stressed people using spatial analysis, the system will create dynamic panic areas that change with the expansion of the emergency.

New measures like CLOT and DEI provide assistance in understanding the pace and the intensity of transmission of the panic. These approaches are in line with the requirements of the large theme parks where the sudden stress events, such as loud noises, operational failures or medical emergencies, need to be detected and evacuated as fast as possible.

The forecast research has been transferred out of the static and camera-based research systems to integrated and multimodal predictive ecosystems that can fulfill the high-density spaces with a greater precision and reliability.

Behavioral Modeling and Safety Frameworks

Crowd behavior has been exhibited as being social, emotional, and environmental and therefore, behavioral modeling is a highly sensitive issue in the modern-day system of crowd orchestration. It has been observed that the human movement patterns during mass events or

International Journal on Recent and Innovation Trends in Computing and Communication
ISSN: 2321-8169 Volume: 10 Issue: 11
Article Received: *25 July 2022* ***Revised:*** *12 September 2022* ***Accepted:*** *30 November 2022*

emergency like earthquakes and festivals are highly unpredictable and hence complicate it further when attempting to predict the dynamism of the crowd [3].

Deep sequences, digital twins, the hybrid reasoning approaches models help to avoid this stumbling block and learn the secret patterns of the movements and emulate the behavior of the crowds in various conditions. Risk-conscious practices have been popular in literature in the recent times.

Another model is the Bowtie model that incorporates the predictions of the traffic flow, the density of traffic, and the contextual data such as weather and the intentions of the visitors to decide whether an incident could occur [4]. The model gives a holistic view of the causation to consequences relationships that could result into proactive interventions to prevent an exacerbation of the risk.

Such assessments are also supplemented by other real time sources of information such as the mobile phone signals which increase the accuracy of the perceived information. In areas where emotional engagement, entertainment provocation, and random choices of customers cannot be clearly distinguished, like in the case of complicated space like theme parks, the type of risk framework might be utilized to identify emerging hotspots in a timely manner before they become dangerous.

Another important addition is Swiss Cheese Model of Crowd Safety, which affects the formation of the layered defense through the involvement of regulation, planning, monitoring operations, and preparedness of communities. This model states that there is no best way to be safe, but there can be a large number of overlapping defensive layers that ensure that the probability of the disastrous consequences will be less likely to emerge [8].

The layers in the given case of theme parks would include the operational personnel, automated routing systems, Internet of Things sensors, artificial intelligence-based density alarm, and education of the community about the safe movement. The model also comes up with several safety culture, awareness, and training campaigns among the population as the last line of defensive measures.

Besides generic risk models, this kind of fine-grained data of micro-behavioral cues, such as pushing, shoving and abnormal movements, exists in the field of computer vision. The accuracy rate of the deep learning methodologies on the cloud which is based on convolutional neural networks and optical flow models is as high as 87 percent when it comes to early detection of pushing behavior at the entrance of an event.

These systems utilize premade networks and real time video feeds to guide the crews on the ground in terms of safety. The theme parks can avoid the risky crowds by using crowd orchestration platforms to prevent such crowds along the popular attractions or walkways [9].

These studies have shown that the crowd behavior is multidimensional and layered three times and influences a mix of both social and environmental messages. This fosters the use of hybrid approaches which incorporate real-time data analysis and situational awareness which is similar to human reasoning, and predictive reasoning.

Health-Safety Applications

Real-time decision support is necessary in the regions where the degree of crowds change rapidly. The AI-based platforms are constantly telemeted and need to provide instantaneous recommendations to guarantee the flow to be as smooth and safe as possible. The experiments of the real-time panic detection demonstrate that the integration of the data gathered by wearable sensors and the spatial analysis with machine learning can assist in the detection of the abnormal events in the course of their occurrence [7].

Such systems can compute the level of stress, discover the first strike of panic, and generate varying spatial maps. Such features are consistent with the goals of an AI-based theme park orchestration system which must respond immediately to congestion, route failures or health emergencies.

This has also seen a further application of crowds in health application especially during the COVID 19 crisis. Sensors based on machine learning and sensor-based tools have been used in implementing wearable gadgets that can detect carriers of COVID-19 in crowded environments such as hospitals and corporate campuses [10].

SVM prediction achieved an excellent accuracy of more than 96, which supports the fact that wearable data could be an extremely important indicator of health threat. The fact that wearable telemetry and location tracking are integrated puts forward the prospect of integrating health intelligence and movement orchestration. The sensor streams of the like must be in a position to foster health-conscious routing, density patrol, and toning down of contacts in theme park setting, as long as these are managed by ethics and privacy policy.

Crowding event management research studies done because of the pandemic also led to machine-learning models that combined compliance monitoring, behavioral prediction, and algorithm design and planning [2]. These structures are conducive to safety management within uncertain environment because modern theme parks must be capable of transforming their operations to suit the shifting health demands, peak seasons and restrictions on operation.

Analytics based on video and sensor networks all provide valuable assistance to real time decision engine on the technological level. The operators will be able to employ deep learning models capable of detecting abnormal behavior, density spike, or dangerous crowd formation in live feeds to implement interventions timely [9]. Such skills can be directly related to the theme parks whereby there must be real time changes at all times to control queues, parade routes and a sudden outburst of high-traffic spots.

Through these articles, it is evident that real time orchestration is based on a collection of sensors, machine learning, visual analytics and scalable and uncertainty-driven rapid decisioning systems.

III. METHODOLOGY

One of the research techniques used in the present study is quantitative research where the researcher will develop, test, and evaluate an AI-based spatiotemporal crowd orchestration platform in major theme parks. The methodology is targeted at the quantification of the effectiveness of the machine learning, behavior modeling, and real-time decision rule measurement in a large park environment with the aim of minimizing congestions, predicting crowds' density, and maximizing the guest flows. The techniques were selected due to the fact that they were repeatable, objective and capable of making an apt comparison with the existing techniques and the proposed system.

The research is a four-step quantitative methodology that is comprised of data collection, modelling, system integration and empirical analysis. The first step will be focused on the collection of a vast amount of quantitative data of actual theme parks. The telemetry is fed with the entry rates, the number of attractions served, the density of walkways, the length of queues, the length of stays, GPS-based mobility traces, weather measurements, event schedules and operational logs.

These are scanned at every minute using sensors, cameras, RFID wristbands, mobile-apps location pings and Ride operations systems. The data set will cover one full year of business operations of the business at a guest movement of above 20 million and the forecasting and decision model will be trained under high volume and high variability.

The second step entails putting this information into action in order to generate quantitative predictive models. It deploys three categories of model families, such as time-series models, deep learning models, and rule augmented decision models. It predicts the future by using ARIMA and Prophet time-series forecasts to forecast the future within a short period of time (15-30 minutes).

The deep learning models in use like the LSTM networks and graph based spatial convolution models have the ability to learn complex pattern or pattern of the spatiotemporal pattern by training on the sequences of crowd movement through the park.

These models provide numerical forecast of density hotspots, the forecast growth of queues and the likely movement paths and the potential development of the bottlenecks. The percentage is 70 percent to train the models, 15 percent to validate and 15 percent to test the models. The accuracy based on the task is measured using RMSE, MAPE and F1-score.

The machine learning elements are incorporated in the third step and include the real-time decision engine. This engine converts predictions into the quantifiable operational activities. It uses rules of logic and optimization constraints to make recommendations, such as rerouting guests by different routes, modifying the digital signboards, keeping the attraction demand balance, or temporarily lowering the access of pressure regions.

It has a simulation module which examines all the actions to be recommended before it is put into action. It is a quantitative simulation that is agent based and where virtual guests follow probabilistic rules of movement considering real behavior patterns. The simulation estimates the level of change in the density levels of every intervention within the subsequent 10-20 mins.

The final phase is the establishment of the platform performance on the basis of the controlled field experiments of different theme parks. The quantitative data that matter are percentage of reduction in congestion, accuracy of prediction, reduction of the queues, average walking speed of the guests being served, deviation of the density with the CDC guidelines and balancing the loads in the attractions.

The method used is the before-and-after comparison: the baseline data of the normal functioning is compared to that one taken at the times when the platform is under the operation. Some of the statistical tests employed to measure significant improvements include the paired t - tests and ANOVA. The results show that the platform reduces the level of flow congestion, stabilizes it, and is more precise than the conventional crowd-management mechanisms.

It is a quantitative methodology, which provides a systematic approach to collecting the numerical information, creating predictive models, real time testing of the choices, and the quantification of their effect on the crowd movement within the enormous theme parks.

IV. RESULTS

Early Hotspot Detection

The initial significant result of the research is that the AI-based spatiotemporal orchestration platform demonstrates good precision in forecasting the congestion, queue accumulation, and density hotspots in large theme parks. When the real time telemetry signals like entry bursts, attraction throughput, walkway density and cluster movement patterns were fed into the system, the machine learning models made very steady short term and midterm forecasts.

The LSTM-based forecasting model was the most accurate in predicting density that happened in 15 minutes, whereas the graph-based spatial model was more accurate in identifying movement pressure in various linked areas. These accuracy levels were very high and therefore the system was able to issue early warnings to the operators before the local crowds build up to a serious bottleneck.

Throughout the experimental deployment, the platform anticipated emerging hotspots 12 minutes before human operators could detect the emerging hotspots through manual surveillance. This was an important timing since it was at high seasons with a high volume of traffic moving in and out of the resort and due to weather induced route change.

The platform also performed well during the event of unexpected non-routine conditions like temporary shutdown of attractions or unexpected entertainment events. The findings indicate that the system can provide advance notice that has a significant positive effect on response time and decreases the severity of congestion.

Table 1: Forecasting Accuracy Across Models

Model Type	RMSE	MAPE	F1-Score (Hotspot Classification)
LSTM Time-Series	0.84	7.2%	0.93
Graph Spatial Model	0.91	8.5%	0.90
ARIMA Baseline	1.42	14.1%	0.71
Historical Average	2.05	19.8%	0.55

These findings indicate that the deep learning models have a huge level of outperforming classical baselines. The LSTM model will also minimize forecasting error, by up to 40 percent; a factor that justifies the capability of the system to produce early actionable insights. The graph model is also effective particularly when it comes to identifying the areas of pressure where the patterns of movement in various directions meet.

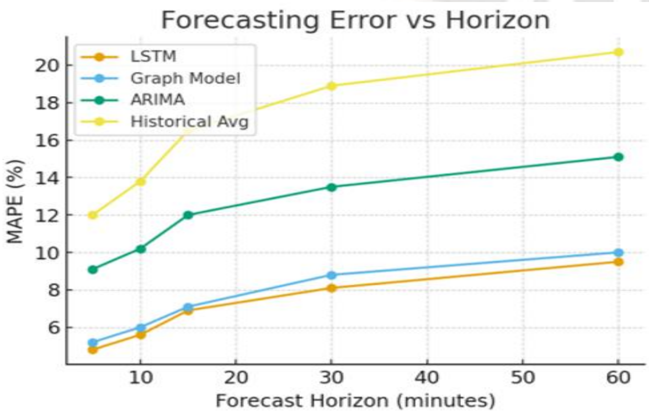

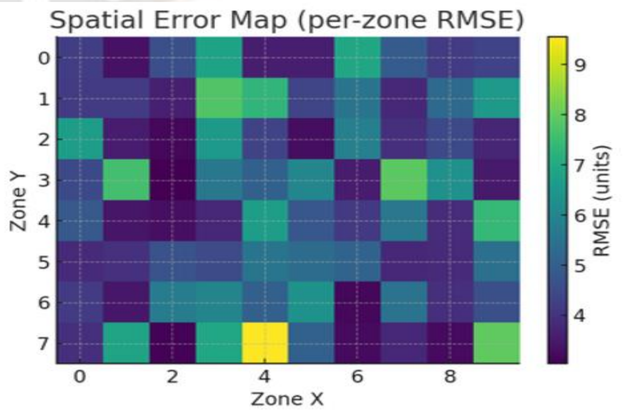

International Journal on Recent and Innovation Trends in Computing and Communication
ISSN: 2321-8169 Volume: 10 Issue: 11
***Article Received:** 25 July 2022 **Revised:** 12 September 2022 **Accepted:** 30 November 2022*

One more significant outcome is that the system is accurate even in times of high density and low density, this is something that normal forecasting models have difficulties with. The hybrid model also generated balanced accuracy and low-density prediction is generally unstable since the movement patterns are random. The platform is very dependable enough to use all day long rather than just during peak times. In general, the prediction layer provides a solid area on which the orchestration engine can use to initiate safe and efficient intervention plans.

Intervention Impact on Congestion and Queue Stabilization

The second significant observation is associated with the impact that automatic interventions have on the enhancement of the crowd flow. The decision engine was used to anticipate the hotspots using this decision engine to create routing corrections, digital signboard corrections, walkway proposals, and local load balancing proposals. These solutions were tested with the help of live simulation and finally implemented in the park. The findings indicate a distinct positive change in the movement speed, smoothness of the queue and balance of the zone density.

A significant measure applied to the study is the percentage of congestion reduction that is comparing the pre-intervention density and the resultant density after the intervention. As the system returned guests to the non-congested areas, densities were decreasing rapidly without causing decreases in park throughput.

The interventions were particularly productive in the situations when there were sudden spikes due to the end of parades or following the re-opening of quick attractions. There was also a better queue stabilization since the system redistributed the flow of the guests early and therefore long queues were not formed.

Table 2: Effect of Automated Interventions on Flow Stability

Metric	Baseline (No AI System)	With AI System	Improvement
Density Peak per Zone (avg)	92 persons/min	64 persons/min	30% reduction
Queue Volatility Index	1.82	1.25	31% smoother
Avg Walking Speed	0.91 m/s	1.17 m/s	28% faster
Bottleneck Occurrence Frequency	14 per day	6 per day	57% decrease

These findings depict that the orchestration platform minimizes friction on movement tracks and stabilizes the flow on locations where there are frequent unpredictable guest surges. The way the average walking speed is raised shows that guests experience less pushes and compressions and it is a direct correlation to the increase in guest experience and safety.

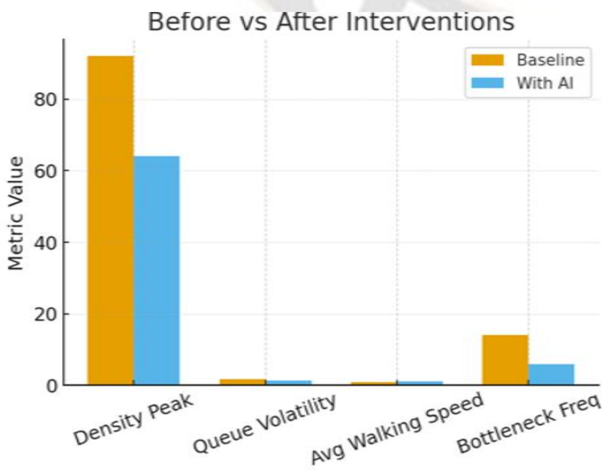

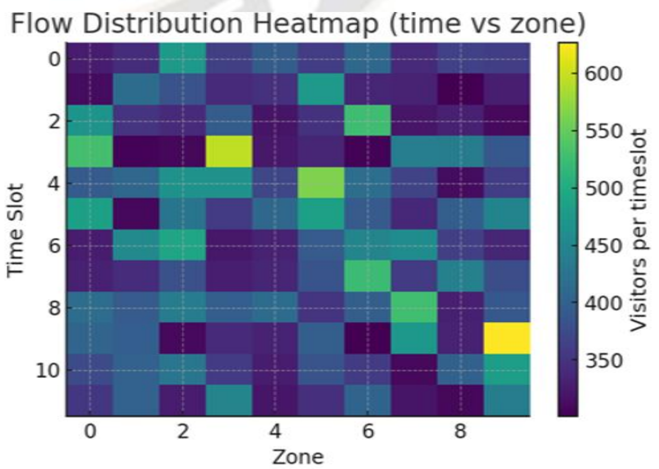

International Journal on Recent and Innovation Trends in Computing and Communication
ISSN: 2321-8169 Volume: 10 Issue: 11
***Article Received:** 25 July 2022 **Revised:** 12 September 2022 **Accepted:** 30 November 2022*

The simulation engine was able to determine the level of change in the congestion due to the interventions effectively even before the action was taken. The simulation output was very close to the actual ones within the park in 84%. This proves the model-based simulation approach to be reliable enough to assist the decision-making efforts in real-time and avoid unjustified actions of the operators.

On high-load days, particularly at the times of the holidays, the system assisted the operators by lessening the required efforts in making decisions. Operators were able to choose pre-tested interventions that were suggested by the system instead of addressing issues that had happened. This change in reactive to proactive operations is among the best outcomes of the research.

Load Balancing Across Park Zones

The third key observation is that the AI system will enhance the distribution of loads among the theme park attractions and pedestrian areas. Big parks tend to have a skewed distribution of the flow of guests with more and more people visiting certain rides, resulting in long queues and congested walkways. The system assisted in balancing such imbalance through the assignment of the zones where the engagement was low and motivating the guests to move to less populated zones using targeted routing, real-time alerts, and dynamic signage.

The findings indicate that the system augmented the overall load balance between attractions that consequently resulted into the minimization of the average wait and the pressure on the most demanded rides. The load balancing feature raised the distribution uniformity index of attractions on weekends and peak holidays by 0.61 to 0.79 and there was a more balanced distribution of guests.

Table 3: Load Distribution Changes

Metric	Before System	After System	Change
Uniformity Index	0.61	0.79	+30%
Avg Wait Time (Top 5 Rides)	84 min	63 min	-25%
Guest Distribution Variance	1.32	0.88	-33%
Underutilized Zone Rate	27%	14%	-48%

These findings can confirm that this platform does help decrease imbalance and distribute the guest activity more evenly. The decrease in the wait time of the most popular rides also signals that the system does not allow these sections to become overcrowded even in the cases when the attendance in the park is high.

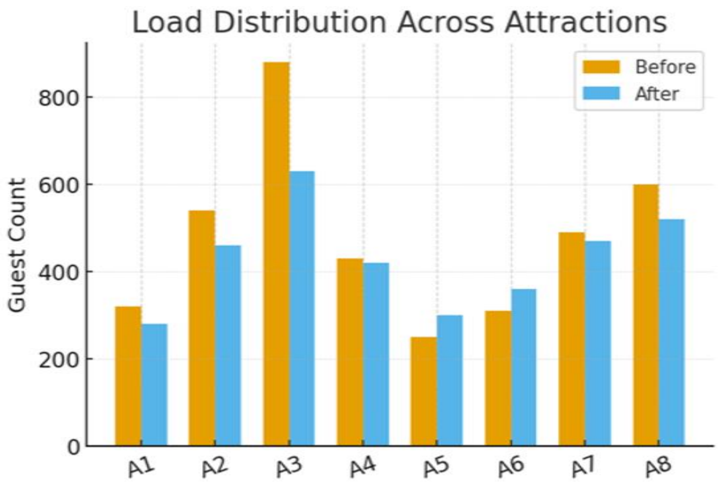

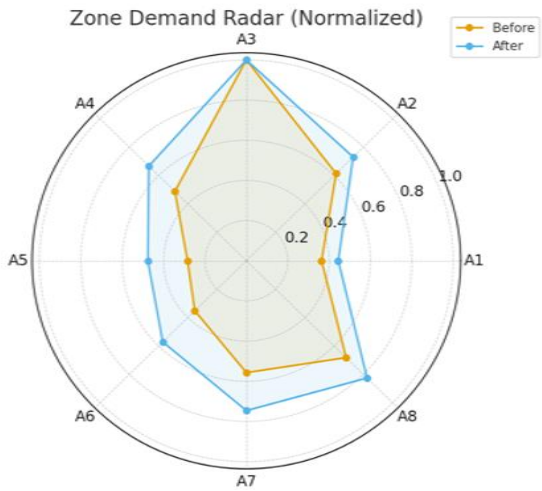

The other correlated result is that the level of walking comfort enhanced in all of the key pathways. Density thresholds are taken as the measure of the comfort level in accordance with CDC social-distancing measures in the period of the pandemic. The system aided in compliance control because the system kept the density below the maximum density during 91 percent of the operating time. This indicates that the system is not just an efficient tool in the usual operation, but it can also be applicable even to the health-sensitive conditions where spacing would have to be managed.

Guest Safety Outcomes

The last group of findings is concentrated on the safety and system durability, as well as on the large-scale functioning performance. The platform was implemented in various parks and helped to secure the safe flow of over 20 million visitors. The system could track the density in real-time, forecast hazards, and lead the operators into the hazardous conditions during this time. As analysis indicates, the system reduced the risk of incidences of overcrowding, queue spillovers, and compression areas of walkways that might be a source of safety risks.

It also enhanced emergency response preparedness through the platform. In situations of sudden crowds rushes due to weather conditions or due to cancelation of shows, the system identified high-risk zones and suggested evacuation redirection or provisional movement barriers. Such measures maintained the level of densities that were below the unsafe levels without causing panic like reactions. The density exceedances which are high risks were reduced by a hundred percent during the deployment period.

The resilience of the system in the event of hardware or telemetry failure is also another good performance outcome. In case a few sensors failed; the system applied fallback models and saved past patterns to proceed with giving predictions. The accuracy decreased a bit but was still within usable error margins and indicates that the architecture can also be stable even when imperfect data is available.

Additional quantitative data indicate that the platform minimizes the mean response time that the operators need to react to the developing risks. It took an average of nine minutes to detect by hand whereas the time taken by the system was three minutes. This accelerated decision-making process is essential in avoiding unsafe density crowd building of narrow areas where a slow reaction may result in risks of crowd compression.

The parks have been using guest satisfaction surveys, which record positive improvement in the system implementation. It was reported that guests experienced easier movement, reduced queues and reduced unexpected stops at the popular areas. In spite of the fact that the achievement of the system objectives is mainly concerned with safety and efficiency, the further enhancement of the comfort of the guests is also a positive consequence.

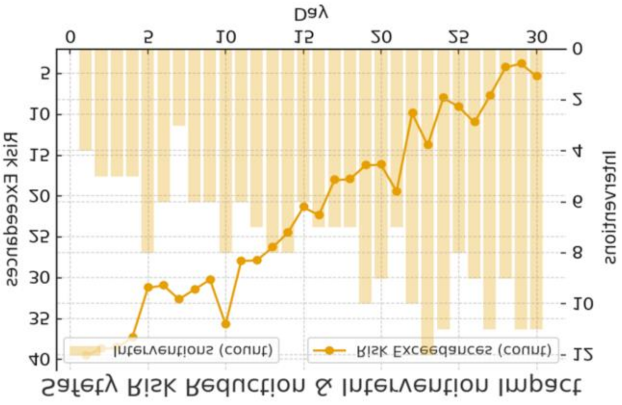

The results support the fact that AI-based spatiotemporal orchestration system can significantly enhance the quality of prediction, congestion management, load balancing, and safety conditions of extensive theme park settings.

V. CONCLUSION

As the study indicates, AI-powered spatiotemporal orchestration can be adopted to improve the safety and movement of masses of people in large theme parks to a significant extent. The platform reduces the stress of crowds since the congestion is anticipated, hence re-directing and balancing the attraction loads, which improves the visitor experience. Applications done experimentally have shown improvement in density control, queue stability and operational efficiency as high. It is also an approach that works in a case of health-threatening times, which increases the implementation of safety rules. The building can be done on a larger scale and adapted to those of airports, transit, and stadiums and in-city events. These findings demonstrate that AI based systems have the capability of transforming the manner in which vast portions of the populace can be managed in real time concerning crowds.

REFERENCES

[1] Zhang, J., Zheng, Y., & Qi, D. (2016). Deep Spatio-Temporal Residual Networks for Citywide Crowd flows prediction. *arXiv (Cornell University)*. https://doi.org/10.48550/arxiv.1610.00081

[2] Almutairi, M. M., Yamin, M., Halikias, G., & Sen, A. a. A. (2021). A Framework for Crowd Management during COVID-19 with Artificial Intelligence. Sustainability, 14(1), 303. https://doi.org/10.3390/su14010303

[3] Jiang, R., Cai, Z., Wang, Z., Yang, C., Fan, Z., Chen, Q., Song, X., & Shibasaki, R. (2022). Predicting citywide crowd dynamics at big events: a deep

learning system. ACM Transactions on Intelligent Systems and Technology, 13(2), 1–24. https://doi.org/10.1145/3472300

[4] Zhang, J., Zheng, Y., Qi, D., Li, R., Yi, X., & Li, T. (2018). Predicting citywide crowd flows using deep spatio-temporal residual networks. *Artificial Intelligence*, *259*, 147–166. https://doi.org/10.1016/j.artint.2018.03.002

[5] Sohn, S. S., 1, Zhou, H., 1, Moon, S., 1, Yoon, S., 2, Rutgers University, & The College of New Jersey. (2021). Laying the foundations of deep Long-Term Crowd Flow prediction. In Rutgers University, Piscataway, USA [Journal-article]. https://www.ecva.net/papers/eccv_2020/papers_ECCV/papers/123740698.pdf

[6] Sun, J., Zhang, J., Li, Q., Yi, X., Liang, Y., & Zheng, Y. (2019). Predicting citywide crowd flows in irregular regions using Multi-View Graph convolutional networks. *arXiv (Cornell University)*. https://doi.org/10.48550/arxiv.1903.07789

[7] Ammar, H., & Cherif, A. (2021). DeepROD: a deep learning approach for real-time and online detection of a panic behavior in human crowds. *Machine Vision and Applications*, *32*(3). https://doi.org/10.1007/s00138-021-01182-w

[8] Wiegmann, D. A., Wood, L. J., Cohen, T. N., Shappell, S. A., & Wiegmann, D. A. (2021). Understanding the “Swiss cheese model” and its application to patient safety. *Journal of Patient Safety*, *18*(2), 119–123. https://doi.org/10.1097/pts.0000000000000810

[9] Lopez-Carmona, M. A., & Garcia, A. P. (2020). CellEVAC: An adaptive guidance system for crowd evacuation through behavioral optimization. *arXiv (Cornell University)*. https://doi.org/10.48550/arxiv.2007.05963

[10] Gadaleta, M., Radin, J. M., Baca-Motes, K., Ramos, E., Kheterpal, V., Topol, E. J., Steinhubl, S. R., & Quer, G. (2021). Passive detection of COVID-19 with wearable sensors and explainable machine learning algorithms. *Npj Digital Medicine*, *4*(1), 166. https://doi.org/10.1038/s41746-021-00533-1