PATH FINDER: An XAI system that enables wayfinding in social contexts

B. M. Elema
Eindhoven University of technology
Industrial Design
1832239
b.m.elema@student.tue.nl

G. J. Vogels
Eindhoven University of technology
Industrial Design
1577514
g.j.vogels@student.tue.nl

J. M. Bosma
Eindhoven University of technology
Industrial Design
1432621
j.m.bosma@student.tue.nl

N. Kamel
Eindhoven University of technology
Industrial Design
1828126
n.kamel@student.tue.nl

ABSTRACT
This project presents Path Finder, an explainable AI based navigation system designed to improve movement and comfort in crowded festival environments. Combining YOLOv8n-based crowd detection with Dijkstra's algorithm for path optimization, the system guides users along less congested routes while visually explaining its reasoning. A prototype was made and the testing demonstrated technical feasibility but also revealed limitations such as static data and a lack of real-world validation. The results highlight the potential of combining AI-based perception, a path finding algorithm and explainability to enhance both navigation efficiency and user comfort in dense social settings.

Keywords
Machine Learning, crowding, navigation, smart navigation systems, Explainable Artificial Intelligence (XAI)

1. INTRODUCTION

1.1 The particular problem that you are going to address.
Crowding is an important aspect affecting human emotions and experiences (Zhang et al, 2023). From an environmental psychology perspective, Stokols (1973) defined the term crowding perception as an individual's subjective value judgement, meaning it is an individual's psychological perception of actually or potentially experiencing inadequate space (Zhang et al, 2023). Crowded and dense environments can lead to confusion and/or irritation, leading to poor experiences of an environment. Research has shown that crowding plays a significant role in the destination choices of individuals, since it can lead to negative emotions associated with crowding, which may affect the user's intent to revisit destinations (Hadinejad et al, 2022) (Yin et al, 2019). Furthermore, the Arousal Theory depicts that, in the context of social crowding, crowding can subsequently instigate negative emotional responses like feelings of stress and anxiety. This is because individuals require a certain extent of spatial, social and psychological safety distances, and these negative emotions are aroused when strangers invade this personal space (Stemmer et al, 2022).

Crowd control is an inherently complex problem that involves numerous factors that need to be taken into account, such as crowd size and density, available routes, an individual's goals, crowd dynamics, psychological and physical condition of individuals, etc. Other challenges are the potential of panic or the vulnerabilities of certain users with specific needs, such as elderly individuals, people with disabilities or simply people who get anxious in crowdy spaces. All these factors not only increase the opportunity for effective management to ensure safety, but also real-time path recommendations for effective and safe movement (Alafif et al, 2025).

Examples of disasters where crowding was a problem are the Astroworld Festival in Houston (2021) and the Love Parade in Germany (2010). Even though surges were the major cause behind the fatalities in these events, bottlenecks and overcrowding near stages or exits are also common (Ticket Fairy, 2025).

The main design objective of the project discussed in this paper is: how can an intelligent and user-friendly crowd control system and application be designed, to guide visitors to their desired location on a festival through the least crowded path? To address this objective, we present PathFinder, a mobile application that is connected to a system with a camera that uses computer vision and a pathfinding algorithm to find the path of least resistance. Within this application, the user can select the spot they aim to go to on the festival, and a compass will lead them while visually showing and explaining to the user why the path is most beneficial, using Explainable AI. Explainable Artificial Intelligence (XAI) aims at producing Machine Learning models with good interpretability for the users and makes the results and decisions made by the model comprehensible to humans, avoiding a "Black Box" approach (Ali et al, 2023). This way, the user experiences autonomy in the process of navigation, but at the same time the organizers of these events are able to exert crown control and possibly preventing escalated crowding situations. While our solution is targeted for all visitors of crowdy events, we aim the application especially for vulnerable groups like people with specific needs or disabilities, elderly or people who suffer from anxiety in crowdy areas, since it is recognized that their need for effective crowd and safety management is more crucial (Alafif et al, 2025).

This paper first presents the related works, followed by methodology explaining the learning algorithm and XAI, the design and implementation of PathFinder as well as the testing and analysis. Lastly, some conclusions, future works and limitations are discussed.

1.2 Related work

Crowd dynamics and crowd analysis

Crowds dynamics have been researched in several domains, including the sociological domain, the behavioral and psychological domain, "level of services" and the computer vision domain (Soraia, 2010). In behavior psychology, Fruin (1971) has examined how crowd density and personal space can affect the crowd dynamics and an individual's comfort within crowdy spaces. Fruin presented a model, the level of services for the design of public spaces (Fruin, 1971). The model presents the conditions and characteristics of people in a crowd for different levels of crowd densities, flow, and movement of the crowd. Hall's (1959) research highlights the social use of the place and proposed the term "Proxemics". It depicts the significance of maintaining interpersonal spaces as a comfort zone which can be reduced in crowded environments. Research into crowd dynamics (Soraia, 2010) has also revealed different behaviors that lead to crowd formations because people tend to take the path of least effort. Examples are the Bottleneck Effect and the Lane Formation (Soraia, 2010). The Bottleneck Effect is when people change velocity due to restrictions in the environment and the density of a crowd, for example near a small exit. The Lane Formation emerges since it is more effortless for people to follow behind others who are already moving in their direction than pushing their own way through the crowd (Soraia, 2010).

Crowd Computer vision

Crowd analysis within computer vision focuses on automated analysis methods for crowded scenes. There are three main models developed for people counting and density estimation models. The first, pixel-level analysis, is focused on density estimation rather than people counting since low-level local features are used (e.g. background subtraction or edge detection) (Soraia, 2010). Texture-based analysis relies on extracting patterns from images for object recognition, using statistical and structural texture descriptors in order to characterize contrasts, coarseness and directionality. Lastly, object-level analysis (Soraia, 2010) aims to detect individual objects in a scene and are more accurate, however, only in low-density crowds. In high-density crowds, this method struggles with occlusions. Yang et al (2003) proposes multi-camera systems to segment people and estimate the number of people in crowds to improve the accuracy. Ali & Shah (2008) propose an approach that goes beyond counting and actual real-time crowd management, meaning crowd tracking and understanding behavior to identify patterns of movement and predict congestions.

Crowd computer vision with Deep Learning (DL)

In more recent developments in computer science, research has also looked into modern deep learning (DL) approaches (Alasmari et al, 2024). Tripathi et al (2019) highlights convolutional neural networks (CNNs) can be used for crowd behavior analysis as well as enhancing tasks for crowd counting, crowd density estimation and crowded scene analysis. Li et al (2014) described crowd analysis in terms of the Dynamics Concept. This models crowds either as a set of individuals (Social Force Model) or as continuous fluids. Kiran et al (2018) looked into the detection or prediction of anomalies, meaning detecting unseen objects or defining rare events (Alasmari et al, 2024). This was done using DL, unsupervised and semi-supervised methods.

Other studies reveal more holistic approaches to understanding crowd behavior, detection and prediction using DL. Al-Shaery et al (2020) provides a review of crowd management systems with decision support systems (DSS) and systems that have early warning systems for crowd safety. Ebrahimpour et al (2019) delved into more holistic crowd analytics by looking at crowd social media analysis (exploiting check-in from geo-tagged blogs on social media), crowd spatiotemporal analysis (uses data from transportation that has GPS like shared bikes and buses) and crowd video analysis.

However, most of these advancements focus on isolated tasks like counting or detection, but do not integrate these specializations into a real-time decision-maker for users. This was the opportunity that the concept of PathFinder is based on.

Smart crowd navigation applications and XAI

Concurrent with these advancements, smart crowd navigation systems that are integrated into mobile applications have emerged. These often use GPS and sensor data to help users navigate from one place to another while providing real-time information about which crowded zones or congested routes to avoid within urban and indoor spaces. Examples include Google Maps (2025), Waze (2025), Crowdless (2025) and CityMapper (2025). What is missing in these applications, is that they often serve as "Black boxes", which means they provide little to no information on why a certain route is recommended. However, it is the transparency and explainability which enhances the user's trust in the recommendations, which is seen as an opportunity in this paper.

Overcrowding and vulnerable groups

Studies have shown that in general, overcrowding negatively impacts mental health. The term "Urban anxiety" describes the mental challenges that people face in densely populated urban areas. It usually manifests as stress, anxiety, sensory overstimulation due to overcrowded spaces and constant noise. Moreover, factors like congestion and the absence of personal space increase these feelings which causes individuals to feel disconnected and overwhelmed (Karwa and Chauhan, 2025). However, especially vulnerable groups are disadvantaged by overcrowding. For example, research has shown that overcrowding has been identified as the most significant factor related to poor mental health for older adults since it restricts physical and emotional space (Firdaus, 2017). Even though these findings were found within an indoor (housing) context, it shows a strong possibility that overcrowding in an outdoor space could still negatively impact elderly. Moreover, people with intellectual disabilities or people with autism are more prone to developing anxiety symptoms due to overcrowding (Al-Jabri et al, 2023). Therefore, there is an opportunity for an application with explainable navigation for the target groups that are especially sensitive to crowding, helping them to take calmer routes that are less congested.

All in all, previous studies have mainly explored monitoring of crowds using AI and XAI in decision systems. However, a clear design gap is the combination of these elements into a system accompanied by a user-centered application hat aids with navigation in real-time. The concept of PathFinder fills this design space by combining crowd analysis with a user-friendly navigation application that stimulates the user to make certain choices based on explained and vizualized reasoning.

2. Methodology & Design Solution

2.1 Design of the interaction

To ensure a user-friendly interface, it was decided to make a mobile app where users can select what festival they are at and where they would like to go. In the app, users will see a map with their current location and different popular locations such as the bars, food

trucks and the stage. Once the user has chosen the location they want to go to, the app analyzes the crowd to find the shortest path from the user's location. The app will then display an arrow pointing in the direction where the user has to go (see figure 1). On the same page a map of the festival is also visible, with an overlayed heatmap and the full route the user will take to give them a better understanding of why the app is making certain decisions. While walking the route, the arrow will keep updating to show you where to go and eventually lead you to your destination.

A key feature of the user interface is the implementation of XAI, by showing the user a heatmap of the crowd, with their current location as well as the recommended path to their destination overlayed. By doing this, the user gets a better understanding of why the app is recommending them a small detour. Together with this is also shows how much faster/less resistance the path is expected to be compared to the Euclidian distance. Finally, the app has an information screen informing the user of how the density of the crowd is analyzed and that no faces or other forms of personal data will be recorded.

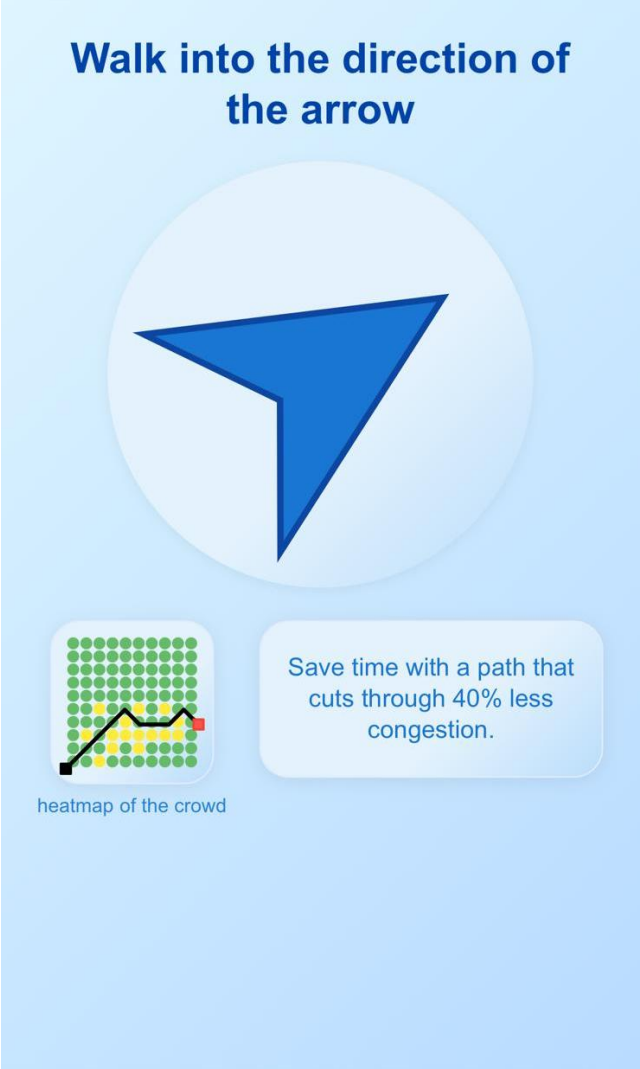


Figure 1: The app guiding the user along the path of least resistance

2.2 Implementation

To realize our concept, the back-end system was divided into multiple parts. The backend consisted of 1) computer vision for crowd detection and 2) a pathfinding algorithm, and the frontend consisted of a web based application to resemble an app that shows the information to the user.

Computer vision

To collect the relevant crowd data, the YOLOv8n computer vison model was used (Ultralytics, nd). YOLO (You Only Look Once) was chosen due to its high speed image processing and lightweight computational requirements, allowing the system to detect individuals in real time with sufficient accuracy. This also made it possible to use a RaspberryPi with a camera module and maintain relatively high frames per second while running the program, compared to more complex models. YOLO detects objects (in our case individuals) by drawing bounding boxes around them (see figure 2&3) (Raspberry Pi, nd). The centre points of these bounding boxes are then mapped onto a 2D grid, where each cell represents X and Y coordinates corresponding to a specific area on the festival floor (see figure 2&3). The value of each cell shows the number of individuals detected in that area, creating a real-time representation of local crowd density. To get accurate results on a 2D space even though the camera module was not situated directly above the environment but rather on an angle, a homography transformation was applied. The information of each cell was stored in a live 2D array, which was used as input for the pathfinding algorithm.

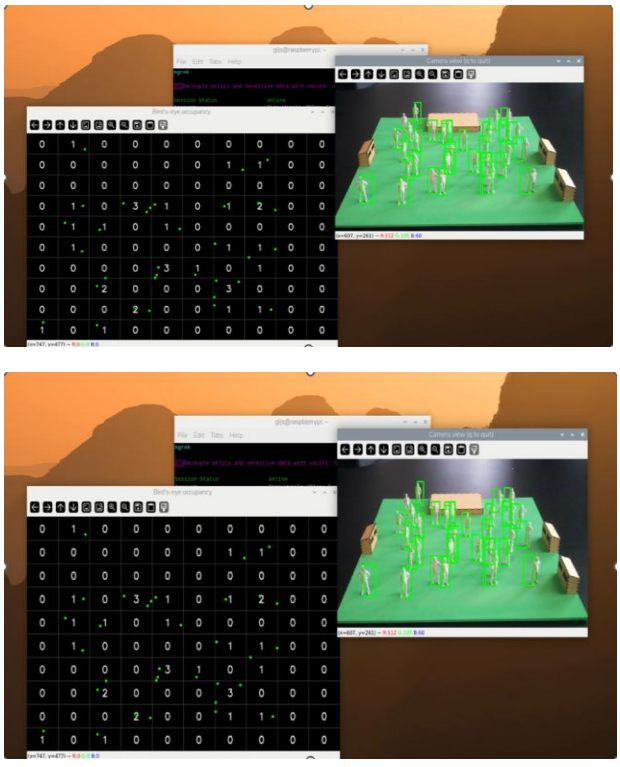


Figure 2&3: YOLOv8n detecting and counting small figures on the scale model

Pathfinding algorithm

To calculate the path of least resistance, multiple options were explored. These options included reinforcement learning and mathematical models such as Dijkstra's algorithm and its alternatives like A* (GeeksforGeeks, 2025). While reinforcement learning offers the advantage of considering multiple dynamic factors, such as crowd movement patterns and temporal changes, it also requires significantly more training time and computational requirements. While using reinforcement learning would be more appropriate in future versions that take into account dynamic crowd movements, Dijkstra's algorithm was used for the prototype since it was deployed in a static environment, which will be explained in the following section. Even though A* would be even faster than Dijkstra's algorithm, Dijkstra was chosen to be the better option as it is more accurate, and processing times were still negligible. Overall, it was able to efficiently compute the path of least resistance from time to time, balancing between travel distance and crowd density while minimizing computation time.

Web-application

To simulate how the system would be deployed in a real world setting, the collected data and path was communicated to a web app, simulating the UI of an app on a user's phone. To communicate between the Raspberry Pi and the web app, Ngrok was used to create a secure HTTP tunnel to transfer data in real time between the Pi's local server and the online web app (Ngrok, nd). The infomation from the Pi was sent as JSON data through Ngrok's public URL. In this way, users could choose where they want to go on the festival layout, after which the interface showed the path of least resistance and provided a compass for navigation.

2.3 Testing and Analysis

For testing and demonstrating purposes, a physical scale model was made. It was chosen to use a physical model over digital crowd generation because it would make a more clear demonstrations. Even though digital crowd generation would allow for easier altering of the heatmap, the physical model would also require a functional implementation of the YOLOv8n model for counting people, making it as close as possible to a real life situation.

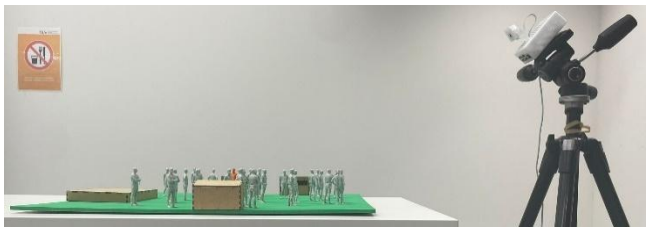


Figure 4: The raspberry pi and camera pointed towards the scale model

To test the efficiency of the YOLOv8n, a confusion matrix was used. Here the number of true positives (TP), false positives (FP) and false negatives (FN) were compared. The number of true negatives are in this case irrelevant as there are no people it should not detect.

Using these numbers, the precision, recall and F1 scores were calculated. Precision is the accuracy of positive predictions. Recall is the sensitivity of the model, indicating how many of the available people it correctly counted. The F1 score combines the Precision and Recall and is the overall score of the model.

$$Precision = \frac{TP}{TP + FP}$$

$$Recall = \frac{TP}{TP + FN}$$

$$F1\ score = 2 \times \frac{Precision \times Recall}{Precision + Recall}$$

Average Precision	94,16%
Average Recall	88,28%
Average F1-score	90,79%

After running the test 4 times, the scores above were calculated. Overall the F1-score of over 90% is reasonable, though it could have been better. This is mainly because of the low Recall score, indicating that just 88% of people present were actually detected. The most probable cause for this is that the YOLOv8n model is meant for detecting humans where in this case it had to detect small, 3D-Printed figures which are less detailed than actual humans. Therefore, it is likely that when testing the model in a real world scenario the Recall score and F1-score will become higher. The only added challenge here would be the people being more densely packed.

3. Discussion

As highlighted in the introduction, crowding has been shown to negatively affect emotions and spatial comfort (Zhang et al., 2023; Stemmer et al., 2022). Therefore, the main goal of this project was to explore the potential of designing an intelligent navigation system that enables users to move through crowded festival environments more comfortably by minimizing their exposure to high-density areas. Reflecting on the prototype implementation and its outcomes, insights were gained on how to design for this goal on both a conceptual and technical level. In this section, we further evaluate the results in relation to the initial objectives and relevant literature, discuss the main limitations of the project, and outline directions for future work.

Starting with the YOLOv8n computer vision component, the model performed reliably in identifying individuals and estimating local crowd density within simulated festival environments. When converted into a matrix, this data provided a manageable yet informative representation of the spatial layout. Although simplified, this approach proved sufficient to visualize crowd distribution and effectively support the routing process. Moreover, the results suggest that overhead cameras positioned above the crowd could serve as a reliable means of tracking crowd distribution, offering valuable insights for the development of future real-time navigation systems. Overall, the implementation of Yolov8n in our project reflects the evolution described in the related works on computer vision, where deep learning models have increasingly been applied for crowd counting and density estimation. However, Path Finder extends these efforts by using vision-based input not only for analysis but for real-time decision-making

The pathfinding module, based on Dijkstra's algorithm, proved accurate in generating routes that avoided highly crowded zones without significantly increasing travel time. This outcome reflects one of the core goals of the system: achieving a balance between route efficiency and route comfort. Throughout the project, it became clear that reinforcement learning and crowd prediction models could further enhance the system's performance in real-world settings by anticipating dynamic crowd movements (an aspect discussed more deeply in the limitations and future work). Within the static and simulated festival environment used for the prototype (where reinforcement learning would have been a technological overkill that could only slow down the system), however, Dijkstra's algorithm proved to be a stable and sufficient solution.

Throughout the app's design process, the importance of XAI was consistently recognized. This project serves as another example where a clear integration of XAI principles can support the achievement of the system's overall goals. In the case of Path Finder, explainable features such as the heatmap, route visualization, and directional arrow were integrated to inform users about the data available to the system and the reasoning behind its decisions. According to the studies and theories discussed in the related work, such transparency is expected to enhance the interpretability for the users and make the results and decisions made by the model comprehensible, letting the user experience autonomy in the process of navigation.

Reflecting on the project, it becomes clear that AI algorithms hold substantial potential for improving crowd navigation at large-scale events. By linking technological innovation with insights from environmental psychology and user-centered design, the project bridges the gap between theoretical studies on crowding behavior and practical, AI-driven crowd management solutions. This report underlines both the technical viability and the social relevance of AI-assisted navigation in crowded settings. Nevertheless, several important limitations of this project also need to be addressed.

3.1 Limitations

First, the detection performance of the YOLOv8n model was evaluated only four times during this project. Conducting more extensive testing and recalculating metrics such as the F1 score would provide a more reliable assessment of its accuracy. Moreover, the results are likely to vary when the system is applied to real-world crowds. In the prototype, crowd density was simulated by placing small 3D-printed figures to represent people. However, since YOLOv8n was not trained on such artificial models, its detection accuracy might improve when applied to real humans. It could, however, also be decreasing in high-density crowds where overlapping bodies could obscure visibility and reduce detection accuracy. In addition, environmental factors such as low lighting, rain, smoke, or the presence of large structures like stages or tents could further challenge the system's ability to maintain reliable crowd-density estimates across the entire festival area.

A second major limitation lies in the static nature of the crowd data used in our project. In the current system, Dijkstra's algorithm processes the crowd distribution only once from one initial snapshot, meaning that crowd flow dynamics are not incorporated. This limits the model's ability to predict and adapt to sudden crowd movements. Consequently, while the system performs well in a static setting, its real-time responsiveness remains untested. Hereby, Dijkstra's algorithm is effective in static environments but less suitable for continuously evolving conditions. Although reinforcement learning was identified as a promising alternative capable of handling dynamic data, it was not fully implemented or evaluated within the project's timeframe. As a result, the performance of a true machine learning–based routing seems promising, but remains uncertain.

Another important limitation concerns user experience and usability. Since the app was not tested with real festival visitors, its overall usability remains unknown. Furthermore, the XAI elements, such as the heatmap and directional indicators, were designed to make the system's reasoning more transparent, but it is still unclear whether these features really enhance user understanding. It is possible that too much visual information could overwhelm users or even lead them to ignore the system's recommendations in favor of their own judgment. In addition, the actual impact of the system on reducing perceived crowding or stress has not yet been empirically validated.

3.2 Future works

Building on both the promising results of this project and it's limitations, several directions for future work will be proposed to enhance both the technical performance and the user experience of a navigation system in the festival setting.

Firstly, large-scale testing of the person detection component is essential. While YOLOv8n performed well in our set-up, its real-world accuracy remains to be validated under the complex and variable conditions of actual festivals. Future studies should examine how detection reliability changes when monitoring large crowds with diverse appearances, movements, and lighting conditions. Such testing will help identify the model's limitations and guide possible adaptations or fine-tuning for outdoor, high-density environments. In addition, it could be valuable to compare YOLOv8n's performance with other computer vision models, as YOLO's speed advantage may come at the cost of reduced accuracy in detecting visitors.

Secondly, the system should evolve from static to dynamic crowd analysis. Implementing real-time updates of crowd movement would allow continuous route recalculation instead of relying on single-frame snapshots. In addition, predictive modeling could be introduced to anticipate the future flow of people. A model capable of forecasting crowd trajectories would make the system proactive, enabling it to suggest alternative routes before congestion occurs. Reinforcement learning appears particularly promising in this regard, as it can adapt over time by learning from continuously changing crowd data. To ensure optimal performance, future research should also focus on training and testing the routing algorithm using real-world metrics such as average route length, average crowd exposure, route duration, and user satisfaction. These evaluations would allow fine-tuning of the algorithm's weighting parameters to better reflect the system's ultimate objectives: improving navigation efficiency, safety, and overall user experience in crowded environments.

Another important direction involves training and testing the model across different types of events. Each event exhibits unique

behavioral patterns (e.g. variations in dancing intensity, walking speed, or spatial usage) that influence how crowds form and move. By incorporating event-specific data, the system could become more robust and context-aware, thereby improving its generalizability beyond a single festival setting.

Finally, further development and evaluation of the interaction design are needed to ensure that information is presented effectively to visitors. A prototype like our app should be tested on a larger scale to assess usability, clarity, and its impact on perceived safety and navigation confidence. As discussed in the related works, there are multiple approaches to delivering route information to users, and future research should explore and compare these alternatives to identify the most effective methods. Regardless of the chosen approach, we argue that explainable AI will remain essential, ensuring that users not only follow the system's guidance but also understand and trust its decisions.

4. Conclusions

The goal of this project was to create an AI-assisted navigation system that allows individuals to move through crowded events such as festivals more safely and easily. The system aimed to translate crowd density data into navigation guidance, allowing visitors of these events to move to important spaces through the path of least resistance. To realize this, a computer vision algorithm was combined with a pathfinder algorithm, after which the gathered data was linked to a web app, simulating a user's phone.

To evaluate the concept in a controlled manner, the system was deployed within a static miniature environment, consisting of a scaled-down festival layout with 3D printed miniature human models to represent the crowd. This setup made it clear to see how the system responded to different crowd locations and densities. Using the YOLO object detection model, the system was able to recognize the miniatures relatively successfully, with some minor false positives and negatives. The number of detections were mapped on a 2D grid using homography, where each cell represented a specific location within the environment, after which the collected data was stored in a live array and used as input for the navigation algorithm.

Since the testing environment was static, a Dijkstra algorithm was used instead of reinforcement learning to calculate the path of least resistance. This algorithm provided efficient route calculations based on distance and crowd density without excessive computational requirements. The data was visualized in a webapp that showed the 2D grid with density information, the shortest route, and a compass to represent the direction the user should walk in.

The results demonstrate that combining computer vision with a pathfinding algorithm can create an intelligent real-time crowd navigation system. Future work should focus on implementing reinforcement learning to evaluate dynamic crowd behaviour, allowing the system to adapt to continuous changes during the route. Overall, the project has highlighted how AI-driven tools can have the ability to improve safety and efficiency at crowded events.

5. REFERENCES

[1] Alafif, T., et al. (2025). Toward an integrated intelligent framework for crowd control and management (IICCM). IEEE Access, 13, 58559–58575. https://doi.org/10.1109/ACCESS.2025.3555154

[2] Alasmari, A. M., Farooqi, N. S., & Alotaibi, Y. A. (2024). Recent trends in crowd management using deep learning techniques: A systematic literature review. Journal of Umm Al-Qura University for Engineering and Architecture, 15, 355–383. https://doi.org/10.1007/s43995-024-00071-3

[3] Al-Jabri, B., Alnuwaiser, S., Abdulghaffar, H., Almuhanna, R., Salaam, S., Brika, R., Addas, A., & Bedaiwi, H. (2023). Healthcare experience of pediatric patients with autism spectrum disorders in Saudi Arabia: A cross-sectional study. Pediatric Reports, 15(3), 452–461. https://doi.org/10.3390/pediatric15030042

[4] Ali, S., Abuhmed, T., El-Sappagh, S., Muhammad, K., Alonso-Moral, J. M., Confalonieri, R., Guidotti, R., Del Ser, J., Díaz-Rodríguez, N., & Herrera, F. (2023). Explainable artificial intelligence (XAI): What we know and what is left to attain trustworthy artificial intelligence. *Information Fusion, 99,* 101805. https://doi.org/10.1016/j.inffus.2023.101805

[5] Ali, S., & Shah, M. (2008). Floor fields for tracking in high density crowd scenes. In *Proceedings of the European Conference on Computer Vision* (Vol. II, pp. 1–14).

[6] Al-Shaery, A. M., Alshehri, S. S., Farooqi, N. S., & Khozium, M. O. (2020). In-depth survey to detect, monitor and manage crowd. *IEEE Access, 8,* 209008–209019.

[7] CityMapper. (2025). The ultimate transport app. https://citymapper.com/

[8] Crowdless. (2025). https://crowdlessapp.co/download/

[9] Ebrahimpour, Z., Wan, W., Cervantes, O., Luo, T., & Ullah, H. (2019). Comparison of main approaches for extracting behavior features from crowd flow analysis. *ISPRS International Journal of Geo-Information, 8(10),* 440.

[10] Firdaus, G. (2017). Built environment and health outcomes: Identification of contextual risk factors for mental well-being of older adults. *Ageing International, 42,* 62–77. https://doi.org/10.1007/s12126-016-9276-0

[11] Fruin, J. (1971). *Pedestrian and planning design.* Mobile, AL: Elevator World Inc.

[12] GeeksforGeeks (n.d.) A* Search Algorithm. https://www.geeksforgeeks.org/dsa/a-search-algorithm/

[13] GeeksforGeeks (2025) Find Shortest Paths from Source to all Vertices using Dijkstra's Algorithm. https://www.geeksforgeeks.org/dsa/dijkstras-shortest-path-algorithm-greedy-algo-7/

[14] Google. (2025). Google Maps. https://www.google.com/maps

[15] Hadinejad, A., Gardiner, S., Kralj, A., Scott, N., & Moyle, B. D. (2022). Cognition, metacognition and attitude to tourism destinations: The impact of emotional arousal and source credibility. *Journal of Hospitality and Tourism Management, 51,* 502–511.

[16] Hall, E. T. (1959). *The silent language.* New York: Doubleday.

[17] Junior, J., & Musse, S. (2010). Crowd analysis using computer vision techniques. *IEEE Signal Processing Magazine, 27(5),* 66–77. https://doi.org/10.1109/MSP.2010.937394

[18] Karwa, T., & Chauhan, N. (2025). Urban anxiety: How crowded cities affect well-being and the role of architectural design. *International Journal of Engineering Research, 14.*

[19] Kiran, B. R., Thomas, D. M., & Parakkal, R. (2018). An overview of deep learning based methods for unsupervised and semi-supervised anomaly detection in videos. *Journal of Imaging, 4(2),* 36.

[20] Li, T., Chang, H., Wang, M., Ni, B., Hong, R., & Yan, S. (2014). Crowded scene analysis: A survey. *IEEE Transactions on Circuits and Systems for Video Technology, 25(3),* 367–386.

[21] Ltd, R. P. (n.d.). *Raspberry Pi.* Raspberry Pi. https://www.raspberrypi.com/

[22] Ngrok. (n.d.). API gateway, Kubernetes ingress, webhook gateway. https://ngrok.com/

[23] Stemmer, K., Gjerald, O., & Øgaard, T. (2022). Crowding, emotions, visitor satisfaction and loyalty in a managed visitor attraction. *Leisure Sciences,* 1–23.

[24] Stokols, D., Rall, M., Pinner, B., & Schopler, J. (1973). Physical, social, and personal determinants of the perception of crowding. *Environment and Behavior, 5,* 87.

[25] Ticket Fairy. (2025). Risk register for cultural festivals: From crowding hazards to cultural harm. https://www.ticketfairy.com/blog/risk-register-for-cultural-festivals-from-crowding-hazards-to-cultural-harm#:~:text=1.,dangerously%20(indianexpress.com)

[26] Tripathi, G., Singh, K., & Vishwakarma, D. K. (2019). Convolutional neural networks for crowd behaviour analysis: A survey. *The Visual Computer, 35(5),* 753–776.

[27] Ultralytics. (n.d.). Quickstart. https://docs.ultralytics.com/quickstart

[28] Waze. (2025). https://www.waze.com/nl/apps/

[29] Yang, D. B., Héctor, N., González-Ba, H., & Guibas, L. J. (2003). Counting people in crowds with a real-time network of simple image sensors. In *Proceedings of the IEEE International Conference on Computer Vision* (p. 122).

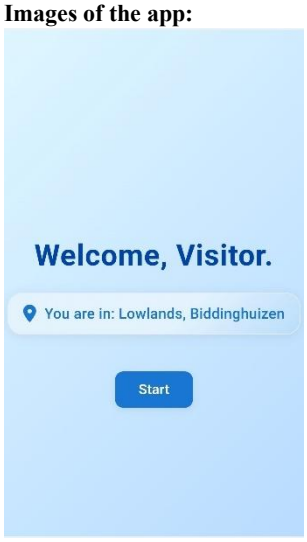
[30] Yin, J., Bi, Y., Zheng, X.-M., & Tsaur, R.-C. (2019). Safety forecasting and early warning of highly aggregated tourist crowds in China. *IEEE Access, 7,* 119026–119040.

[31] Yin, J., Zheng, X.-M., & Tsaur, R.-C. (2019). Occurrence mechanism and coping paths of accidents of highly aggregated tourist crowds based on system dynamics. *PLoS ONE, 14,* e0222389.

[32] Zhang, J., Qi, R., & Zhang, H. (2023). Examining the impact of crowding perception on the generation of negative emotions among users of small urban micro public spaces. *Sustainability, 15(22),* 16104. https://doi.org/10.3390/su152216104

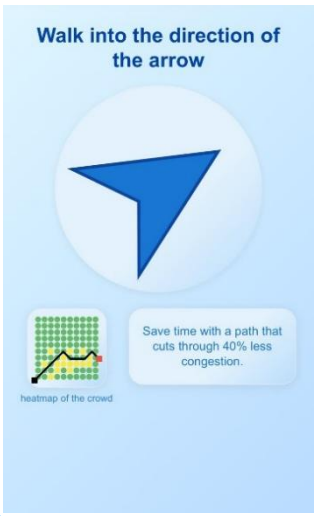
6. Appendix A: Images & Tables

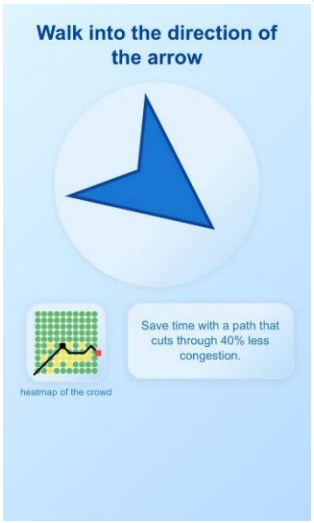
Images of the app:

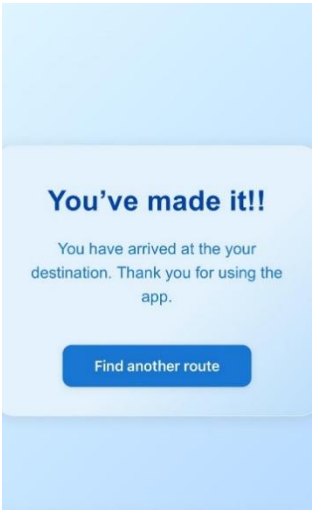








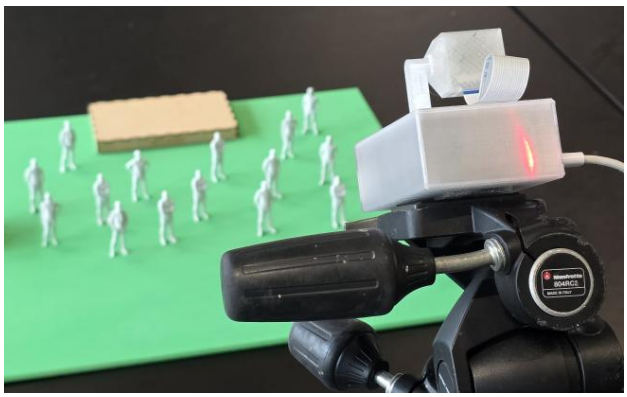




Images of the scale model:







Confusion matrices of the 4 tests done to evaluate the YOLOv8n model:

Total number of people	Number of people correctly counted (TP)
Number of people not counted (FN)	Number of non-people counted (FP)

32	27	32	33	32	23	32	31
5	1	0	2	9	2	1	2

7. Appendix B: Code

Link to github repository:
https://github.com/gijsvogels02/PATHFINDER-DBM140