UAV Remote Sensing and Deep Learning for Automatic Car Detection and Parking Occupancy Analysis: The Case of UANL Stadium, Mexico

Kevin D. Rodríguez González [1], Fabiola D. Yépez-Rincón[1*], Andrea Escobedo Tamez[1], Aylet Vega-Aguilar[1]

[1]Faculty of Civil Engineering, Universidad Autónoma de Nuevo León, San Nicolás de los Garza, Nuevo León, México - krodriguezge@uanl.edu.mx, *fabiola.yepezrn@uanl.edu.mx, andrea.escobedotmz@uanl.edu.mx, aylet.vegaglr@uanl.edu.mx

Keywords: UAV; Football stadium; Car detection; Parking management

Abstract.

Rising traffic demand around university campuses and sports venues exacerbates parking scarcity and congestion. This study develops a UAV–deep learning workflow for the automatic quantification of parked vehicles and the estimation of occupancy levels across facilities at the Universidad Autónoma de Nuevo León (UANL). UAV surveys of the East and West Estadio UANL lots and the FIME faculty lots were conducted with DJI Mavic 2 and Matrice 350 RTK platforms during high-demand periods, including football matches and student egress peaks. The imagery, processed into centimeter-scale orthomosaics (2.4–2.8 cm ground sampling distance), enabled reliable instance detection using a pretrained Mask R-CNN Car Detection model. A total of 4,591 vehicles were identified across the surveyed areas: 2,336 in the West lot, 1,684 in the East lot, and 571 in the FIME lots. Kernel density estimation and spatial metrics revealed near-saturation of stadium lots during matches, reduced occupancy during off-event periods, and elevated but distributed demand in faculty lots during class dismissal. These geospatial indicators were integrated into a parking management framework using heat maps and bottleneck detection around access and egress roads. The approach demonstrates the potential of UAV–deep learning workflows to support demand-responsive parking control, traffic guidance, and long-term planning in congested university and event-driven environments.

1. Introduction

Increasing motorization and concentrated travel demand around universities and sports venues often result in severe parking scarcity and traffic congestion. Parking deficits not only reduce mobility efficiency but also generate negative externalities such as spillover into residential neighborhoods, safety hazards at access points, and elevated emissions from circulation in search of spaces. These challenges are magnified during large events such as football matches, where sudden surges in demand saturate existing infrastructure.

Traditional approaches to monitoring parking utilization manual surveys, ground-based sensors, and closed-circuit cameras face limitations of scale, cost, or field of view. In contrast, aerial remote sensing has emerged as a flexible and scalable solution. High-resolution imagery acquired from unmanned aerial vehicles (UAVs) has proven particularly effective for vehicle detection in urban contexts due to its centimeter-level detail and temporal flexibility (Ammour et al., 2017; Kaya et al., 2023). When coupled with deep learning, UAV imagery enables automatic and reliable identification of individual vehicles, supporting large-scale traffic and parking analyses (Ragab et al., 2023; Yildirim et al., 2024).

Recent developments in convolutional neural networks (CNNs), object detection, and instance-segmentation architectures (e.g., YOLO, Mask R-CNN) have advanced the accuracy and efficiency of vehicle detection in complex urban scenes (Gupta et al., 2021; Srivastava et al., 2021; Bouguettaya et al., 2021). Applications have included car detection in UAV orthomosaics (Ammour et al., 2017; Kaya et al., 2023), traffic speed monitoring in smart cities (Moshayedi et al., 2023), and automated occupancy estimation (Ragab et al., 2023). Nevertheless, empirical studies remain scarce in Latin American urban-university contexts, where parking stress is exacerbated by overlapping academic and sporting activities, limited off-street capacity, and high surrounding urban density.

The UANL, one of the largest public universities in Mexico, exemplifies this situation: its main campus hosts the Estadio UANL football venue alongside major academic faculties, concentrating vehicular flows that frequently overwhelm available parking.

This study develops and implements a UAV–deep learning workflow to automatically quantify parked vehicles and infer occupancy patterns across principal UANL parking facilities, focusing on the Estadio UANL East and West lots and the FIME faculty lots. Specifically, it aims to (i) generate high-resolution UAV orthomosaics of the parking areas, (ii) detect and count vehicles using a pretrained Mask R-CNN model, (iii) analyze occupancy and spatial density patterns under contrasting demand conditions, and (iv) integrate these indicators into a parking management framework to support demand-responsive control and planning.

The contribution of this research is as follows:

- It develops a UAV–deep learning workflow for centimeter-scale vehicle detection and parking occupancy estimation in congested university–stadium contexts.
- It provides the first empirical application of such workflow in Mexico, addressing overlapping academic and sporting demand scenarios.
- It integrates kernel density estimation and spatial analysis into a parking management framework, supporting demand-responsive control, guidance, and planning.

This contribution has been peer-reviewed. The double-blind peer-review was conducted on the basis of the full paper.
https://doi.org/10.5194/isprs-annals-X-3-W3-2025-85-2026 | © Author(s) 2026. CC BY 4.0 License.

1.1 Background

The study site is located at the UANL, in San Nicolás de los Garza, within the Monterrey Metropolitan Area, Mexico. UANL is one of the largest public universities in Latin America, with a student population exceeding 200,000. Its main campus hosts both major academic faculties and the Estadio UANL football venue, which serves as home stadium for Tigres UANL, one of the most popular soccer teams in the country.

The analysis domain covers approximately 34.61 ha, encompassing the East and West parking lots of the Estadio UANL and the parking areas of the Faculty of Mechanical and Electrical Engineering (FIME). This area is characterized by intense fluctuations in parking demand: near-saturation during football matches, high turnover during class hours, and partial vacancy during off-event periods. The combination of academic and sporting functions in a dense urban setting creates a unique environment for testing UAV–deep learning approaches to vehicle detection and parking management.

The regional climate is semi-arid with warm summers, and the study site is embedded in a dense metropolitan fabric, where parking shortages regularly lead to traffic congestion and spillover into surrounding neighborhoods. By focusing on this environment, the research provides insights into mobility challenges faced by universities and sports venues in rapidly growing cities.

2. Study area

The study area is located in San Nicolás de los Garza, within the Monterrey Metropolitan Area (MMA), northeastern Mexico. It comprises the Estadio Universitario (commonly known as *Estadio UANL*), home to the Tigres UANL football club, and its immediate urban surroundings (25°43′21.7″ N, 100°18′ 42.9″ W). The analysis domain (Figure 1) encompasses the stadium bowl and the principal off-street parking facilities that serve both the venue and adjacent university precincts, including the West and East stadium lots and the parking areas of the Faculty of Mechanical and Electrical Engineering (FIME).

Figure 1. Study area

The mapped extent covers approximately 34.61 ha, offering a compact yet heterogeneous environment for vehicle detection and parking-demand estimation. This zone is characterized by pronounced temporal fluctuations in vehicular occupancy: peak saturation during football matches, elevated turnover during university class hours, and partial vacancy during off-event periods. Its mixed academic–sporting function, embedded within a dense metropolitan setting, provides a representative case study of mobility challenges in Latin American cities.

3. Methodology

3.1 Data collection

Figure 2 illustrates the UAV-based workflow applied in this study. Data collection was performed using two platforms: a DJI Mavic 2 Pro and a DJI Matrice 350 RTK. The survey covered ~34.61 ha, including the four principal parking facilities adjacent to the Estadio UANL in San Nicolás de los Garza, Mexico.

The aerial survey used 433 images from an FC2204 camera, achieving a 2.44 cm/pixel GSD; produced 233,281 tie points and 719,111 projections with a mean reprojection error of 0.984 pixels alignment in Agisoft Metashape.

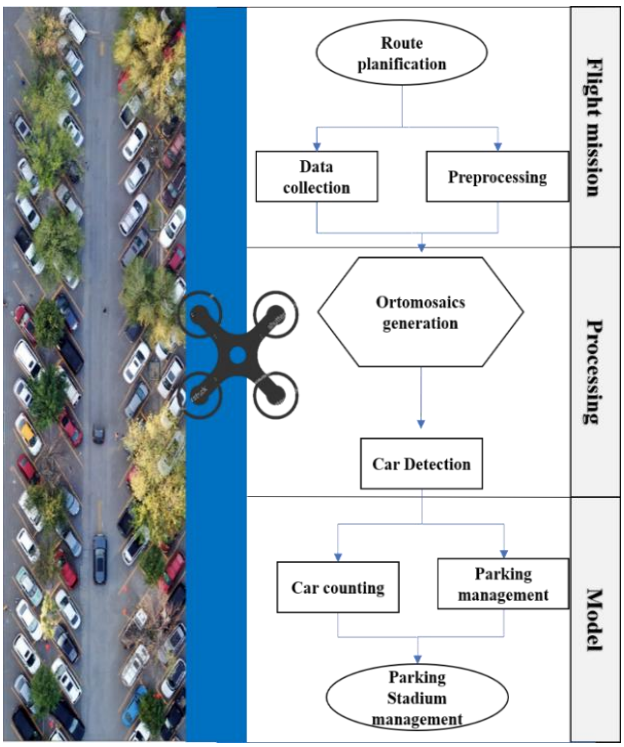

Figure 2. Workflow.

Flights were executed at an average altitude of 82.5 m above ground level (AGL) and a cruising speed of ~8 m/s, producing centimeter-scale ground sampling distance (GSD ≈ 2–3 cm/pixel). Mission planning and execution were carried out in DroneDeploy, employing terrain-following based on a digital elevation model (DEM) to maintain constant altitude and consistent GSD across variable topography.

Standard mission parameters included a lawn-mower flight pattern, 75–85% forward and side overlap, nadir (−90°) camera orientation, automatic exposure, and distance-based triggering. Edge passes were incorporated to ensure complete coverage and strengthen block geometry. Pre-flight checks verified obstacle clearance, GNSS stability, and communications.

The resulting RGB imagery displayed uniform radiometry and geometry, suitable for photogrammetric processing and

This contribution has been peer-reviewed. The double-blind peer-review was conducted on the basis of the full paper.
https://doi.org/10.5194/isprs-annals-X-3-W3-2025-85-2026 | © Author(s) 2026. CC BY 4.0 License.

subsequent vehicle-detection analysis. Complementing the UAV flights, ground-level inspections with handheld digital cameras were conducted to record local reference conditions and support validation of automatic detections.

3.2 Data processing

The UAV imagery was processed in Agisoft Metashape Professional to generate high-resolution, georeferenced orthomosaics of each parking facility. The workflow included automatic internal camera calibration, high-accuracy image alignment, and block adjustment. From this process, dense point clouds and digital surface models (DSMs) were generated as intermediate products, although the primary deliverable for analysis was the RGB orthomosaic.

The orthomosaics preserved the native ground sampling distance of the input images (2.4 – 2.8 cm/pixel) and achieved sub-pixel reprojection error, confirming the geometric stability of the photogrammetric block. Each orthomosaic was exported in WGS84/UTM Zone 14N coordinates and optimized for integration into GIS environments.

Quality control included visual inspection of areal coverage, radiometric continuity, and potential distortions. Minor artifacts were observed in localized areas adjacent to trees or elevated structures, but due to the high redundancy in image overlap, even vehicles partially occluded by canopies remained identifiable in the final products. In cases where inconsistencies were evident, minimal corrections were applied, or the region was flagged as a potential source of uncertainty for subsequent automatic detection.

3.3 Photogrammetry analysis

The photogrammetric survey (Table 1) was designed to generate geometrically stable, centimeter-scale RGB orthomosaics from UAV flights conducted at mean altitudes between 80 and 84 m. This configuration yielded ground sampling distances (GSD) ranging from 2.44 to 2.77 cm/pixel, ensuring sufficient spatial resolution for reliable vehicle discrimination.

The West parking area was surveyed with 195 images across 16.3 ha (GSD = 2.44 cm/pixel), the East lot with 178 images covering 13.1 ha (GSD = 2.60 cm/pixel), and the FIME-P.North sector with 60 images over 5.2 ha (GSD = 2.77 cm/pixel). Minor differences in GSD reflected small variations in effective flight altitude and acquisition geometry, but all missions provided centimeter-level resolution suitable for subsequent deep learning detection.

Processing steps included block adjustment with automatic camera calibration, dense point cloud generation, followed by the production of the high-resolution RGB orthomosaic as the principal analysis product. Sub-pixel reprojection errors confirmed the stability of the photogrammetric block. All deliverables were exported in WGS84/UTM Zone 14N for seamless integration into GIS and deep learning workflows.

Area	Images	Mapping coverage (ha)	Flight altitude	GSD (cm/pixel)
P. East	178	13.1	83.3	2.60
P. West	195	16.3	83.8	2.44
FIME	60	5.2	80.3	2.77

Table 1. Photogrannetry survey

Quality control procedures assessed areal coverage, radiometric continuity, and potential geometric artifacts. In cases where localized inconsistencies were detected such as near wooded areas or elevated structures minimal corrections were applied, or the areas were flagged as potential sources of uncertainty for automatic detection.

3.4 Calibration and overlap

Automatic internal camera calibration was performed during processing. Image residuals revealed minimal deviations from the ideal projection geometry, confirming the stability of the optical system and supporting reliable image orientation.

3.5 Car detection

Car detection is a pre-trained deep learning MaskRCNN model architecture implemented in the ArcGIS API for Python, through the "Detect objects using deep learning" tool (Hou & Li 2024). When local domain adaptation is required, it supports fine-tuning with Train Deep Learning Model. The workflow begins with downloading the model and obtained orthomosaics with an expected spatial resolution of =< 20 cm aproximately (Gui et al 2024).

The detection tool is then configured by specifying the input raster, the output feature class, and the model with *.dlpk* package. Critical inference parameters are set under Model Arguments: *tile_size* (chip size used to tile the image), *padding* (margin to blend adjacent predictions and mitigate edge artifacts, capped at half of *tile_size*), *batch_size* (number of tiles per step, constrained by available memory), *threshold* (confidence cutoff for filtering detections), and *return_bboxes* (Esri. 2021).

Optionally, Non-Maximum Suppression can be applied to suppress lower-confidence overlapping detections, configuring the score field and the maximum allowable overlap. In the Environments pane, the processing extent, a Cell Size of 0.1 (consistent with the working scale), and the processor type are specified; use of a GPU is recommended, when available, to accelerate inference. Upon execution, the model outputs a feature layer containing the detections (class and score), which is automatically added to the map and is ready for post-analysis such as zonal counts, occupancy estimation, or visual validation.

According to the reported accuracy metrics, the Car Detection model in ArcGIS Pro attains an average precision of 0.81 on its evaluation set, indicating strong positive predictive performance with a comparatively low false-positive rate across decision thresholds.

In accordance with the objective of this study, we did not create a local reference dataset with manual annotations, nor did we retrain the pre-trained model. Instead, we adopted an operational validation strategy that comprised: (i) spot visual checks on representative tiles to verify the correct identification of visible vehicles; (ii) a sensitivity analysis of the inference parameters to ensure the stability of the counts; and (iii) spatial consistency checks with respect to observable capacity constraints and expected density gradients near entry/exit points.

Although the workflow is referred to as automated, it should be clarified that the process was implemented as a modular and

automation-ready sequence across interoperable software environments rather than within a single integrated system. The methodology combines UAV image acquisition, photogrammetric processing, and deep-learning-based detection as independent but connected modules through standard geospatial formats (GeoTIFF, shapefile). Each stage can be executed in batch mode or via scripting (Metashape and ArcGIS APIs), ensuring reproducibility and scalability. This modular design enhances adaptability and transparency, prioritizing methodological clarity over closed end-to-end automation.

4. Results and discussion

4.1 UAV photogrammetry

The photogrammetric processing successfully generated high-resolution orthomosaics of the surveyed parking areas, providing a geometrically consistent base for vehicle detection. The workflow began with the alignment of photographs (Figure 3), where common tie points across overlapping images were identified to estimate the position and orientation of each camera station. This stage produced a stable block geometry and confirmed sufficient redundancy for dense reconstruction.

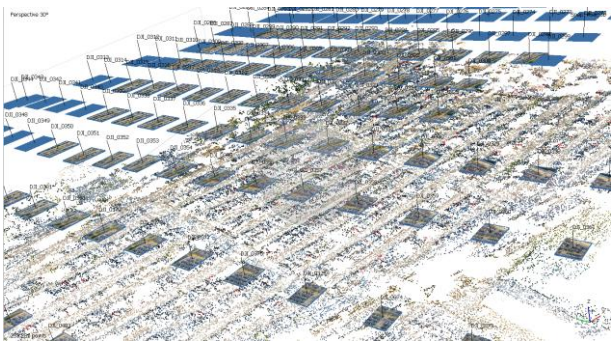

Figure 3. Aligning photos processing

Following alignment, a sparse point cloud was generated to eliminate outliers and non-essential features (Figure 4). This was subsequently densified into a dense point cloud (Figure 5), capturing the structural details of the parking facilities, including rows of vehicles, access lanes, and vegetation. The dense reconstruction was further interpolated into a mesh, which served as the geometric framework for orthomosaic generation.

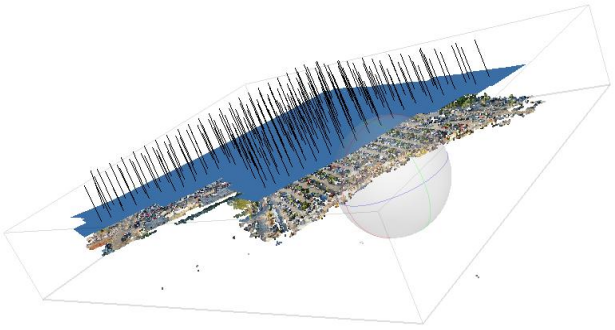

Figure 4. Point cloud with aligned camera positions.

The final orthomosaics corrected for perspective distortions inherent in the original oblique imagery, producing a true orthogonal view with centimeter-level resolution. This enabled reliable visual interpretation of individual vehicles and parking patterns across the 34.1 ha study domain. Comparable levels of geometric precision and radiometric continuity have been reported in other UAV-based vehicle detection studies (Ammour et al., 2017; Kaya et al., 2023; Ragab et al., 2023), supporting the suitability of the dataset for deep learning–based detection.

Figure 5. Dense point cloud, west parking.

4.2 Image orientation

The spatial distribution of camera (Figure 6) stations and footprints indicates continuous coverage and sufficient photogrammetric redundancy, consistent with the number of observations and an average tie-point multiplicity of 3.79, which supports block stability during triangulation.

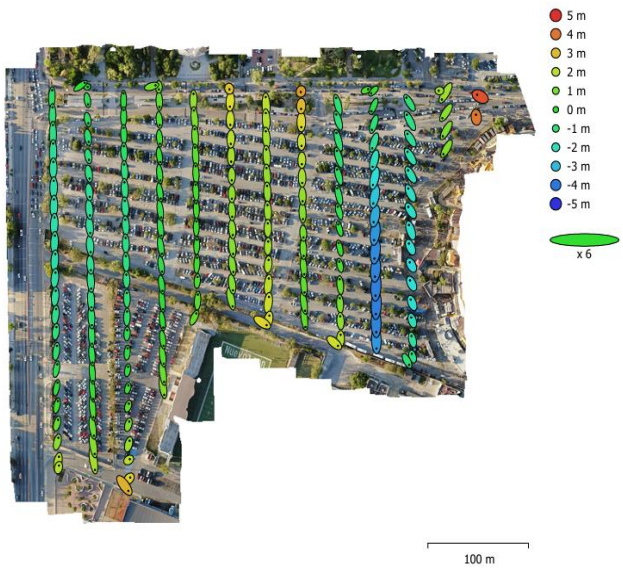

Figure 6. Image orientation

The reported RMS reprojection error ≈ 1.20 px confirms that the estimated poses and calibration parameters coherently explain image measurements at a typical sub- to bi-pixel level for low-altitude surveys with self-calibrated cameras. Collectively, the image count, GSD, projections, and reprojection error validate the geometric quality of the block and the suitability of the overlap pattern for DEM/orthomosaic generation.

Color-coded depth of coverage indicates the number of images per ground pixel (1 to >9). The lawn-mower flight pattern yields banded redundancy; black dots mark camera stations estimated after bundle adjustment. The survey acquired 178 images at 83.3 m (GSD ≈ 2.6 cm/pixel) over 13.1 ha, achieving a mean reprojection error of ~1.2 px

This contribution has been peer-reviewed. The double-blind peer-review was conducted on the basis of the full paper.
https://doi.org/10.5194/isprs-annals-X-3-W3-2025-85-2026 | © Author(s) 2026. CC BY 4.0 License.

High, uniform overlap in the core (≥9 images) ensures robust triangulation and supports reliable ortho/DEM generation, whereas peripheral areas with lower overlap (1–3) are more error-prone and should be trimmed or handled carefully during seamline editing.

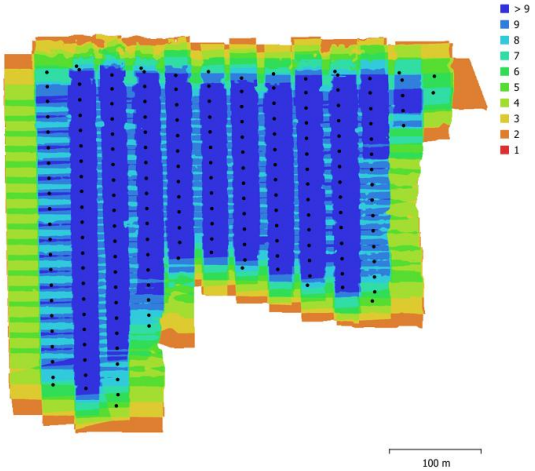

Figure 7. Camera locations and image overlap

The image residuals (Figure 8) summarize the quality of the camera's interior and exterior orientation within the photogrammetric block. Each arrow depicts the residual vector of a tie point—the difference between its measured image location and the position projected by the bundle-adjusted model; the arrow direction indicates local bias, and its length (scaled to 1 px) the error magnitude, typically color-coded (green/yellow: low; red: high). The observed pattern—short, scattered vectors near the image center and longer, partially radial vectors toward the edges—is characteristic of wide-angle optics and denotes residual radial and tangential lens distortion.

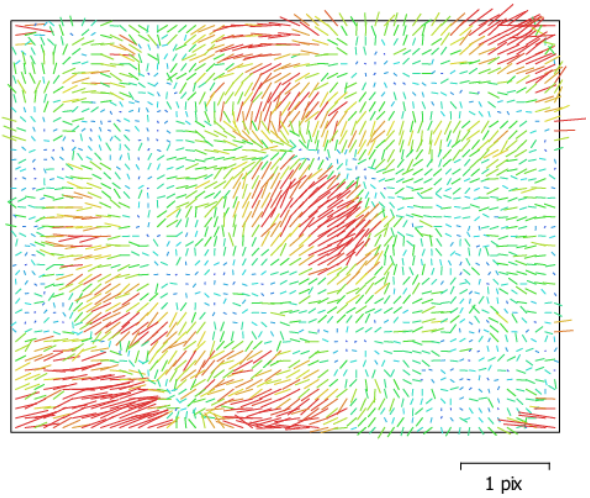

Figure 8. Camera Calibration-Image residual

According to the calibration coefficient and correlation matrix, the interior orientation is consistent with wide-angle UAV optics. The sensor (4000×3000 px; 1.58 µm pixel pitch) implies a physical size of ~6.32×4.74 mm. The adjusted principal distance is F = 3167.76 px, equivalent to ≈ 5.01 mm (3167.76 px × 1.58 µm px^{-1}), slightly exceeding the nominal 4.386 mm as expected from effective focusing and flight conditions. The principal point exhibits modest offsets (Cx = −23.38 px ≈ −36.9 µm; Cy = 11.83 px ≈ 18.7 µm), typical and well bounded. Radial distortion is captured by K1 = −0.0383, K2 = 0.0526, and K3 = −0.0536, indicating an undulating pattern with predominant barrel behavior toward the periphery, whereas tangential distortion is small (P1 = −0.00128; P2 = 0.000737). Reported parameter uncertainties are very low (σ[K1] ≈ 2×10^{-4}; σ[P1] ≈ 4.9×10^{-6}), supporting a stable bundle adjustment and aligning with sub-pixel reprojection residuals.

4.3 Car detection

The parking areas surrounding the UANL stadium display heterogeneous layouts, with the predominant configuration being 45° angled parking, complemented by sections of 90° perpendicular parking, and some areas lacking formal demarcation (Figure 9). In the latter, drivers adaptively combine perpendicular and parallel parking depending on the available space.

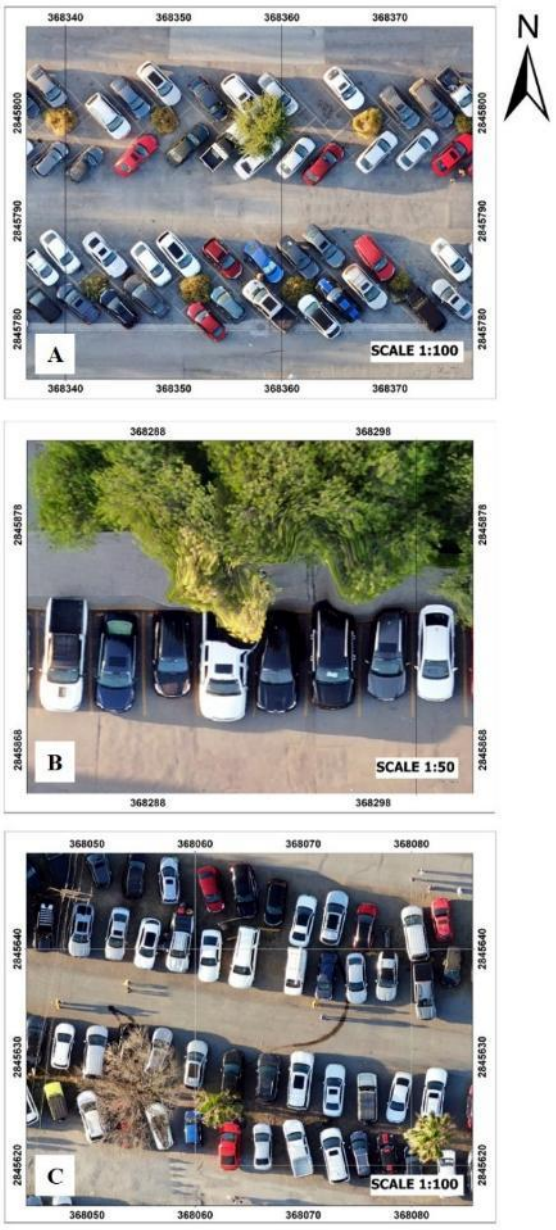

Figure 9. Orthomosaics of the parking lots, showing the different existing layouts. A) 45° angled parking, B) 90° perpendicular parking and C) 90° perpendicular parking without space demarcation.

Using the pretrained Mask R-CNN deep learning model implemented in ArcGIS, a total of 4,591 vehicles were automatically detected across the three analysed parking lots. Table 2 summarizes the vehicle counts per facility, with the West lot containing 2,336 vehicles, the East lot 1,684 vehicles, and the FIME lot 571 vehicles. Detection outputs demonstrated high reliability, consistently identifying vehicles despite challenges such as tree cover or cast shadows. Even when partial occlusions occurred, the algorithm successfully delineated bounding boxes, and each car was correctly registered as a single detection (Figures 10–11).

Parking area	Number of vehicles
P. East	1684
P. West	2336
FIME (P. North)	571

Table 2. Number of vehicles per parking area.

These results are consistent with recent advances in vehicle detection from UAV imagery, where deep learning models have shown resilience to complex urban environments and partial occlusion (Ragab et al., 2023; Yildirim et al., 2024). The detection rate achieved in this study demonstrates a comparable level of precision to other high-resolution UAV applications (Bouguettaya et al., 2021; Kaya et al., 2023), confirming the effectiveness of pretrained architectures for real-world parking scenarios. Unlike prior work focused primarily on traffic monitoring (Moshayedi et al., 2023) or general object recognition (Gupta et al., 2021), this study emphasizes the university–stadium context, a setting characterized by sudden surges in demand linked to scheduled events, which remains underexplored in the literature.

Figure 10. Aerial photo of the stadium and car detection for the entirety of the UANL stadium's parking lots.

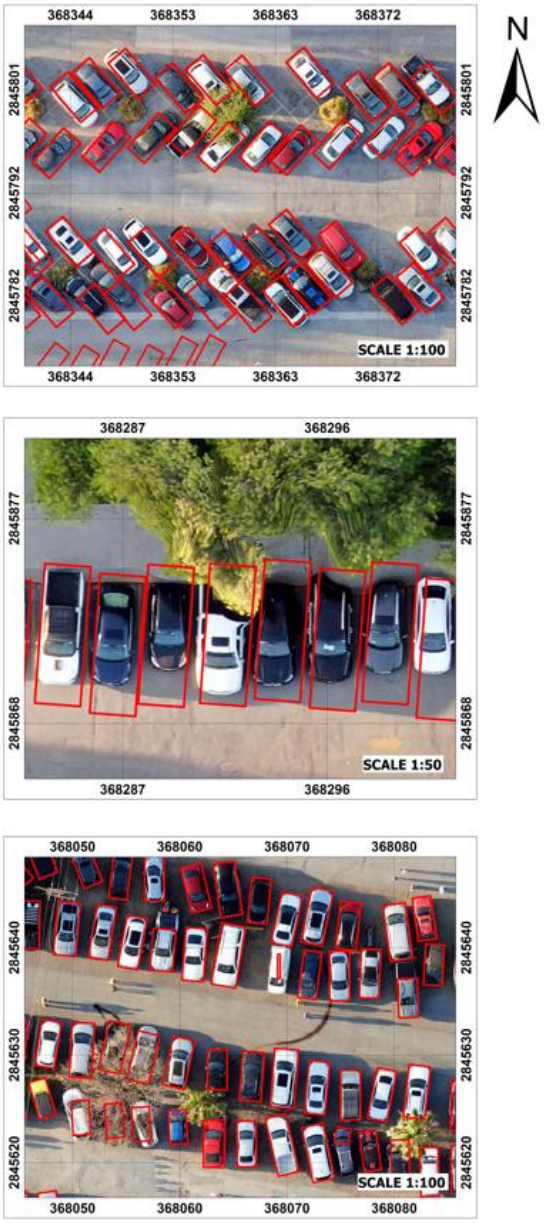

Figure 11. Car detection for the existing parking areas demarcations.

4.4 Density analysis

It is important to note that in this study, the term occupancy is used in a geospatial sense to represent the areal intensity of vehicle presence, rather than the conventional metric of occupied spaces divided by total parking capacity. Because the orthomosaics include both marked and unmarked areas, it was not possible to compute an accurate ratio of vehicles per designated bay. Instead, Kernel Density Estimation (KDE) was employed to transform discrete vehicle detections into a continuous density surface, enabling the identification of congestion hotspots and spatial gradients of use pressure across each facility.

Analysis of detected vehicle counts, and mapped coverage enabled estimation of use pressure for each lot. In Parking East density (Figure 12) with 1,684 vehicles distributed across 13.1

This contribution has been peer-reviewed. The double-blind peer-review was conducted on the basis of the full paper.
https://doi.org/10.5194/isprs-annals-X-3-W3-2025-85-2026 | © Author(s) 2026. CC BY 4.0 License.

ha, vehicular density is high and concentrates around primary access points adjacent to the main avenue.

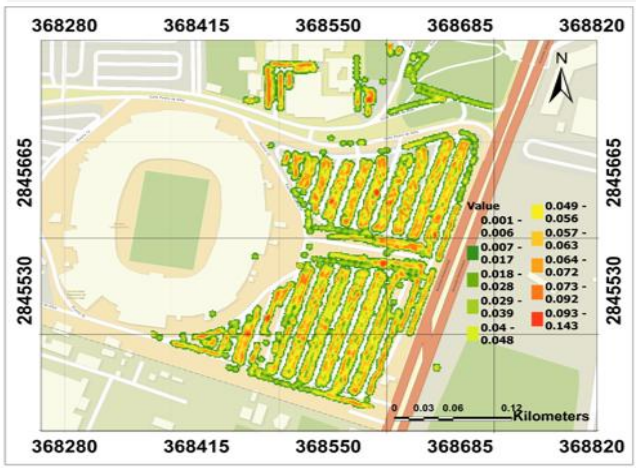

Figure 12. Kernel density P.East

Kernel density estimation delineates critical linear zones foreshadowing bottlenecks at egresses to the arterial network. Under maximum occupancy, clearance times are expected to exceed 25–35 minutes due to concurrent exit maneuvers and the limited capacity of evacuation corridors.

Additionally, in FIME area, with 571 vehicles over 5.2 ha, density (Figure 13) values are less critical; nevertheless, restricted access and adjacency to academic facilities increase vulnerability to internal obstructions and lane blockages. Although relative evacuation capacity is higher, projected peak egress times of 15–20 minutes may affect pedestrian mobility and concurrent academic operations.

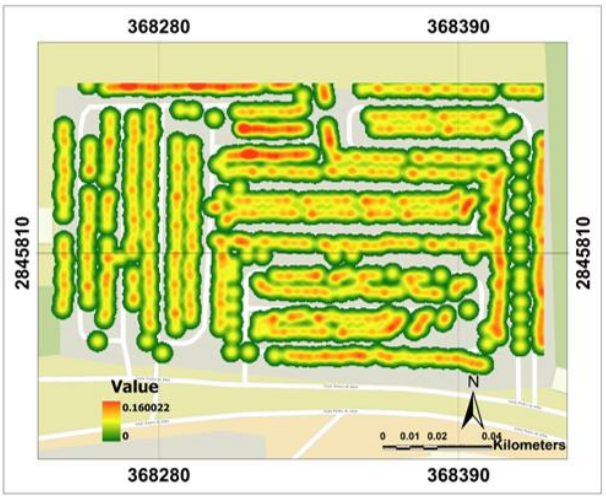

Figure 13. Kernel density FIME

A salient finding is that a portion of vehicles is parked outside designated bays, occupying areas not intended for parking. This irregular (spillover) placement reduces the effective maneuvering envelope within the lots and degrades egress throughput, increasing the likelihood of lane blockages, obstruction of pedestrian access points, and evacuation delays. While exit capacity is relatively adequate under routine conditions, this pattern of informal occupancy can drive clearance times beyond expected ranges and further compromise traffic safety and internal circulation across the university campus

In Parking West, which recorded 2,336 vehicles over 16.3 ha, the greater spatial capacity partially attenuates pressure. Kernel density estimation (Figure 14) reveals a clear spatial gradient in vehicular concentration, with high-congestion clusters rendered in red and, conversely, low-intensity zones depicted in blue to purple, enabling rapid visual discrimination of potential bottlenecks versus relatively uncongested corridors.

However, hotspots identified in the southern sector reveal accumulations that trigger localized congestion. Proximity to the high-demand Avenida Universidad constrains egress by through traffic, with delays estimated at 30-40 minutes during mass departures after sporting events. Despite its size, this lot exhibits a high risk of point saturation capable of propagating congestion onto secondary avenues.

Figure 14. Kernel density P. West

While the total number of vehicles detected in each lot provides the basis for computing exact densities (vehicles/ha), KDE offers additional insight by revealing local concentration patterns that may not coincide with lot boundaries. Future work could integrate detailed parking-space inventories to derive true occupancy ratios and validate congestion metrics under different event scenarios.

4.5 Limitations of the study

Although the workflow demonstrated robust performance in detecting and quantifying vehicles, certain limitations must be acknowledged. First, partial occlusions caused by trees, light poles, and building shadows occasionally led to incomplete bounding boxes, even if vehicles were ultimately counted. Second, the study relied on a pretrained Mask R-CNN model without extensive local fine-tuning; while results were consistent with previous literature (Ragab et al., 2023; Yildirim et al., 2024), domain-specific training could further improve detection accuracy in complex urban scenes. Finally, the analysis was conducted during a single high-demand event, limiting the temporal representativeness of parking dynamics.

Multi-event or longitudinal surveys would provide a more comprehensive understanding of variability across seasons and

event types. Because the workflow purposefully avoids manual annotation to preserve full automation, we did not compute site-specific precision–recall statistics; future extensions may include a small audit sample (≈1–2% of the area) to report local metrics without altering the automated nature of the pipeline.

4.6 Implications for parking and mobility management

The integration of UAV photogrammetry with deep learning offers significant potential to improve parking and mobility planning in high-demand environments such as university campuses and sports venues. The detection of congestion hotspots and irregular parking patterns provides actionable evidence to guide interventions, including redesigning access/egress corridors, allocating staff for traffic control during peak periods, and implementing dynamic guidance systems to direct drivers toward underutilized areas. In addition, the methodology can support the integration of intelligent transportation systems (ITS), combining UAV monitoring with real-time sensors and digital signage to reduce clearance times and enhance safety. Beyond the UANL case study, this approach can be scaled to other urban campuses and metropolitan sports facilities, contributing to more resilient and sustainable urban mobility strategies.

5. Conclusions

This study demonstrated the effectiveness of integrating UAV-derived orthomosaics with pretrained deep learning models in ArcGIS Pro for the automatic detection and quantification of parked vehicles in high-demand environments such as the UANL Stadium. The workflow achieved reliable detection of 4,591 vehicles across 36.4 ha, confirming its suitability for large-scale applications in complex urban and campus contexts.

By combining object detection with kernel density analysis, the study not only quantified parking occupancy but also identified critical congestion hotspots. Results revealed that the East and West lots generate the highest pressure on the external road network, creating recurrent bottlenecks at main exits. In contrast, the FIME lot exhibited lower overall occupancy but showed irregular parking patterns that reduce maneuvering space and increase the risk of internal blockages and evacuation delays.

These findings highlight the potential of UAV deep learning integration as a decision support tool for parking management, providing actionable indicators for mitigating congestion, improving traffic safety, and optimizing evacuation efficiency during mass events. The methodology is transferable to other urban and sporting contexts, contributing to the implementation of intelligent transportation systems and supporting more sustainable urban mobility planning.

Future work should focus on refining detection accuracy under challenging conditions and conducting temporal monitoring across multiple events. Such improvements will enhance the integration of UAV-based monitoring into broader frameworks of smart city governance and campus mobility management.

6. Acknowledgements

Thanks to the Secretariat of Science, Humanities, Technology and Innovation for the scholarship awarded to Kevin Rodríguez CVU:1014573, through the Doctoral Program in Engineering with a focus on Environmental Engineering at the Autonomous University of Nuevo León. Thanks also to the Department of Geomatics and the Faculty of Civil Engineering.

7. References

Ammour, N., Alhichri, H., Bazi, Y., Benjdira, B., Alajlan, N., Zuair, M., 2017: Deep learning approach for car detection in UAV imagery. *Remote Sensing*, 9(4), 312. https://doi.org/10.3390/rs9040312

Bouguettaya, A., Zarzour, H., Kechida, A., Taberkit, A.M., 2021: Vehicle detection from UAV imagery with deep learning: A review. *IEEE Transactions on Neural Networks and Learning Systems*, 33(11), 6047-6067. DOI: 10.1109/TNNLS.2021.3080276

Esri, 2021: Detect Objects Using Deep Learning (Image Analyst). *ArcGIS Pro tool reference*. Accesed 12 September 2025.

Gui, S., son..S., Qin, R., Tang,Y., 2024: Remote Sensing Object Detection in the Deep Learning Era-A Review. *Remote Sensing*, 16(2), p.327.

Gupta, A., Anpalagan, A., Guan, L., Khwaja, A. S., 2021: Deep learning for object detection and scene perception in self-driving cars: Survey, challenges, and open issues. *Array*, 10, 100057. https://doi.org/10.1016/j.array.2021.100057

Hou, T., & Li, J. 2024: Application of mask R-CNN for building detection in UAV remote sensing images. *Heliyon*, 10(19).

Moshayedi, A.J., Roy, A.S., Taravet, A., Liao, L., Wu, J., Gheisari, M., 2023: A secure traffic police remote sensing approach via a deep learning-based low-altitude vehicle speed detector through uavs in smart cites: Algorithm, implementation and evaluation. *Future transportation*, 3(1), 189-209.https://doi.org/10.3390/futuretransp3010012

Kaya, Y., Şenol, H.İ., Yiğit, A.Y., Yakar, M., 2023: Car detection from very high-resolution UAV images using deep learning algorithms. *Photogrammetric Engineering & Remote Sensing*, 89(2), 117-123. https://doi.org/10.14358/PERS.22-00101R2

Ragab, M., Abdushkour, H.A., Khadidos, A.O., Alshareef, A.M., Alyoubi, K.H., Khadidos, A.O., 2023: Improved deep learning-based vehicle detection for urban applications using remote sensing imagery. *Remote Sensing*, 15(19), 4747. https://doi.org/10.3390/rs15194747

Srivastava, S., Narayan, S., Mittal, S., 2021: A survey of deep learning techniques for vehicle detection from UAV images. *Journal of Systems Architecture*, 117, 102152. https://doi.org/10.1016/j.sysarc.2021.102152

Yildirim, E., Sefercik, U.G., Kavzoglu, T., 2024: Automated identification of vehicles in very high-resolution UAV orthomosaics using YOLOv7 deep learning model. *Turkish Journal of Electrical. Engineering and Computer Sciences*, 32(1), 144-165. doi:10.55730/1300-0632.4060