RotaWise: Intelligent Platform for Urban Maintenance

João P. D. da Silva[1], Thiago B. Parente[1], Luan I. F. Santos[1] e Paulo A. L. Rego[2]

[1]Data Insight Wise (DI Wise) – Fortaleza – Ceará – Brasil

[2]Departamento de Computação (DC) – Universidade Federal do Ceará (UFC)

{joaopedro,thiagoparente,luanicaro}@diwise.com.br, paulo@dc.ufc.br

Abstract. *This paper presents RotaWise, an intelligent platform for urban maintenance that integrates cameras and an edge device within a vehicle for real-time video processing and urban issue identification. The system leverages computer vision and AI technologies to monitor roads and sidewalks for problems such as potholes, trash, and debris. The experimental results demonstrated RotaWise's ability to process up to three cameras simultaneously, offering real-time event detection and metadata transmission.*

Resumo. *Este artigo apresenta o RotaWise, uma plataforma inteligente para manutenção urbana que integra câmeras e um dispositivo de borda dentro de um veículo para processamento de vídeo em tempo real e identificação de problemas urbanos. O sistema utiliza tecnologias de visão computacional e IA para monitorar ruas e calçadas em busca de problemas como buracos, lixo e entulho. Os resultados experimentais demonstraram a capacidade do RotaWise de processar até três câmeras simultaneamente, oferecendo detecção de eventos em tempo real e transmissão de metadados.*

1. Introduction

The rapid growth of cities and the increasing population density pose growing challenges to the management of urban infrastructure. In this context, smart cities emerge, applying innovative technologies to optimize public services, improve quality of life, and promote sustainability. Automation in the collection, integration, and interpretation of urban data plays a crucial role, especially in urban maintenance, where the rapid identification and resolution of issues such as potholes, illegal waste disposal, and traffic sign malfunctions are essential for the well-being of the population [Hannan et al. 2015].

Traditional urban monitoring approaches, based on inspections and citizen complaints, present limitations, such as delays in problem identification and resolution, high operational costs, and a reactive management posture when responding to public reports. Therefore, there is a need for technological solutions that enhance the capacity of management teams, facilitating the collection, processing, and real-time analysis of data, enabling rapid actions in the maintenance of urban infrastructure. The advancement of computer vision technologies and edge devices opens new possibilities for automating these processes [Sanchez-Sepulveda et al. 2024].

In this context, we propose **RotaWise**, an intelligent urban maintenance platform that integrates **WiseBox**, an IoT device designed to capture, process, and transmit georeferenced data on the conditions of public roads in real-time. The WiseBox combines advanced hardware and software to identify critical issues such as potholes, debris, and irregular solid waste, and transmit events via 4G/5G networks to the RotaWise backend.

RotaWise aims to support decision-making and optimize urban maintenance operations, with the potential to reduce operational costs, accelerate maintenance, and enhance safety on urban roads. The ability to identify and record issues in real-time strengthens public management, promoting cleaner, safer, and well-maintained cities. This paper outlines the architecture of RotaWise, its technological components, and performance evaluation results conducted with variations in WiseBox configuration and the number of simultaneously processed cameras.

2. Related Work

The advancement of computer vision and edge computing technologies has driven the development of innovative solutions for smart urban management. Several studies explore systems that apply real-time video analysis to monitor urban environments, optimize public maintenance operations, and assist in traffic management. In this section, we present a critical analysis of works related to RotaWise.

The TrafficEd system, proposed in [Chen et al. 2024], utilizes Nvidia Jetson devices for traffic monitoring, detecting illegal parking and counting vehicle flows. While it shares the use of edge devices with RotaWise, TrafficEd focuses on traffic management, whereas RotaWise prioritizes the identification of urban structural issues such as potholes and waste. Similarly, the Polly system [Li et al. 2023] aims to optimize video processing through cross-camera inference, but with a focus on urban traffic surveillance, while RotaWise focuses on urban maintenance.

[Tran et al. 2022] presented a driving assistance system based on mono cameras and Nvidia Jetson devices, while [Kuo and Lin 2024] developed a multi-task perception model for illegal parking detection. Both works use embedded devices and computer vision, similar to RotaWise, but focus on vehicular safety rather than urban maintenance. [Liang et al. 2024] proposed the SplitStream system, which optimizes video processing on edge devices to reduce latency and increase efficiency. While RotaWise also focuses on resource optimization, it differentiates itself by integrating georeferenced data for urban problem mapping, extending the scope of proposed solutions. [Donati et al. 2020] proposed an automated waste detection system for urban cleaning trucks, using computer vision to optimize energy consumption. While both systems aim to optimize urban operations, RotaWise offers a more comprehensive solution for urban infrastructure monitoring.

The authors of [Adi et al. 2024] explored fusion algorithms for parking spot identification using the YOLOv7 model. Although the focus is parking management, the use of object detection models brings this work closer to RotaWise, which also applies AI models for urban object detection. [Wang et al. 2022] developed the EVA system, which optimizes real-time video analysis by suppressing spatial and temporal redundancies, similar to RotaWise's goal of reducing computational resource usage. Other works, such as [Matsuda et al. 2021], presented an illegal parking detection system using dashboard cameras and machine learning algorithms. RotaWise shares real-time analysis but with a specific focus on urban maintenance. [Zhao et al. 2024] proposed the CMCA-YOLO, applying improvements in object detection in environments with occlusions, similar to Rota-

Wise's approach in dealing with challenges in complex urban scenarios. While RotaWise shares the use of Jetson devices, it differentiates itself by focusing on urban mapping.

In summary, Table 1 highlights the main similarities and differences between RotaWise and the reviewed works. The comparison shows that RotaWise stands out by integrating georeferenced data collection with real-time video analysis, offering a comprehensive solution for urban maintenance. This approach enables more complete urban problem monitoring, complementing existing solutions focused on traffic surveillance, vehicular safety, or energy efficiency.

Table 1. Comparison Between Related Works

Paper Title	**Cameras in Vehicle?**	**Processing Location**	**Devices Used**	**Application Context**
[Chen et al. 2024]	No	Edge	Jetson TX2	Traffic Monitoring
[Li et al. 2023]	No	Edge	Jetson Nano	Traffic Surveillance
[Tran et al. 2022]	Yes (Mono)	In the car	Jetson TX2/Xavier NX	Vehicle Assistance
[Kuo and Lin 2024]	Yes	In the car	Jetson Nano	Illegal Parking
[Liang et al. 2024]	No	Edge	Jetson Nano/AGX Xavier	Distributed Processing
[Donati et al. 2020]	Yes	In the car	Jetson TX2	Automated Urban Cleaning
[Wang et al. 2022]	Yes	In the car	Jetson TX2	Vehicle Perception
[Matsuda et al. 2021]	No	In the car	Nvidia GPU	Illegal Parking
[Zhao et al. 2024]	No	Edge	Jetson Nano	Cross Inference
RotaWise	Yes	In the car	Jetson Xavier/Orin	Urban Maintenance

3. RotaWise

The architecture of RotaWise is formed by a set of components that interact to monitor and optimize the conditions of public roads. These components include the vehicle, where the cameras and the WiseBox device are installed, and the cloud service. The vehicle is the mobile agent of the RotaWise platform, responsible for traversing public roads and capturing data through the installed cameras. During its movement through the city, the vehicle collects information about the road conditions, identifying issues such as potholes, trash, and debris. In terms of functionality, the vehicle's main role is to ensure coverage of urban areas and to provide physical support for the computing and data capture devices.

The cameras are responsible for capturing images and videos, positioned strategically to offer a comprehensive view of the road and sidewalk conditions. They generate video streams, which are processed locally by the WiseBox to perform detection and classification of objects of interest. Additionally, the cameras can be configured to automatically adjust their capture settings, such as exposure and brightness, depending on the time of day, thus ensuring better adaptation to lighting conditions and optimizing image quality in different environments.

The WiseBox is the central component of the system, integrating hardware and software for real-time data processing. It is a compact IoT-based device equipped with an Nvidia Jetson module responsible for the local processing of video streams from the cameras. Using computer vision algorithms and AI models optimized with TensorRT, the WiseBox performs tasks such as detection, tracking, and classification of objects on roads and sidewalks. Depending on the Nvidia Jetson model (Nano or Xavier NX), the system can process up to 3 camera streams. Additionally, the WiseBox includes a dual-chip router modem, ensuring stable communication with cloud services and allowing real-time transmission of detected events and their metadata. In cases of connectivity loss, the device stores events locally and synchronizes them when the connection is restored. The WiseBox also includes a GPS module, essential for georeferencing the images, facilitating spatial analysis and decision-making. Figure 1(a) shows an image of the WiseBox and Figure 1(b) illustrates how the cameras are installed in the vehicle.

(a) WiseBox (b) Installed cameras

Figure 1. Images of the WiseBox and cameras installed in the vehicle

The cloud service is responsible for receiving, storing, and processing the data sent by the WiseBox. It manages the events of interest coming from the cameras and the GPS module, presenting these data in a structured way on a dynamic dashboard accessible to the urban management team. Additionally, it is possible to generate reports and metrics that assist in making strategic decisions. Figure 3 shows two images of the dashboard, where it is possible to search, filter, and view georeferenced events on a map, as well as access the images and view their metadata, including the identified objects.

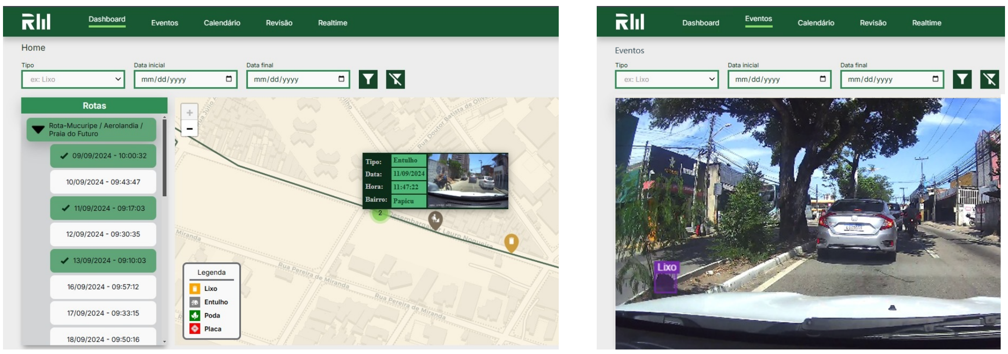

(a) Event visualization on the map (b) Event details

Figure 2. Images of the RotaWise dynamic dashboard

4. Evaluation

To evaluate the performance of the RotaWise platform, controlled experiments were conducted for the detection and tracking of urban waste, such as trash, debris, and tree trimming, on a public road in the city of Fortaleza, in the state of Ceará. Detection was performed using the YOLOv9 model, while tracking was executed by the OCSort algorithm. The detected events were sent to the RotaWise backend in the cloud, where data processing and analysis were performed in real time.

The experimental setup involved the use of different NVIDIA Jetson models on the WiseBox to assess their performance in relation to the number of connected cameras.

The Jetsons Nano and Xavier NX versions were tested, in conjunction with 1, 2, and 3 connected cameras, according to the configuration of each experiment. The choice of the number of cameras was based on preliminary experiments aimed at determining the system's ability to process video streams in real time while maintaining a frame rate of 25 frames per second (fps). In these experiments, it was observed that the Jetson Nano was unable to properly process more than two cameras simultaneously, while the Jetson Xavier NX managed to handle up to three connected cameras, maintaining system performance.

The main concern during the experiments was to monitor the WiseBox's computational resource usage, including CPU, GPU, and memory. These metrics were chosen to assess the processing load imposed by the detection and tracking algorithms, as well as the system's ability to handle multiple cameras in different configurations.

The Figures 3 (a), (b), and (c) show the CPU, GPU, and memory usage of the WiseBox Nvidia device, respectively. As expected, increasing the number of cameras being processed also increases CPU, GPU, and memory usage. The Jetson Nano reaches its CPU and GPU limits with two cameras, averaging 92% usage for both. On the other hand, the Xavier NX can process up to three cameras, reaching a maximum of approximately 82% CPU usage and 72% GPU usage. Since the Jetson Nano is recommended for development (non-production) environments, the results are still noteworthy. With further optimizations, the Xavier NX could support even more cameras, especially when GPU-based video decoding is enabled.

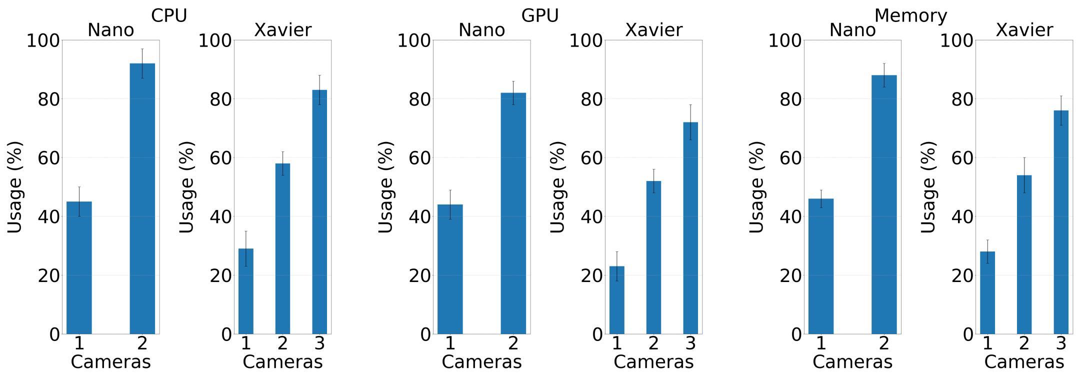

(a) CPU (b) GPU (c) Memory

Figure 3. WiseBox resource usage during the experiments

In summary, the experiments analyze RotaWise's performance under different hardware conditions and amount of camera, allowing for the identification of limitations and the proposition of system improvements. These analyses are also essential for optimizing the application of RotaWise in real-world urban monitoring scenarios.

5. Conclusion

RotaWise demonstrates its potential as an innovative solution for urban maintenance, offering real-time monitoring and intelligent detection of urban issues. Experimental results show its ability to process multiple cameras simultaneously, but there are areas for improvement, particularly in resource optimization. Future work will focus on evaluating the quality of the data collected and the accuracy of the detection results. Additionally, experiments will include scenarios involving connectivity loss, testing the system's ability to store events and resynchronize once the connection is restored. Finally, efforts will be made to improve scalability by supporting more cameras and enhancing detection algorithms to meet the demands of smart city infrastructure. With these improvements, RotaWise will provide a more efficient and scalable solution for urban maintenance, contributing to sustainable development.

References

Adi, G. S., Nugroho, H., Rahmatullah, G. M., Yusuf, M., and Fadhlan, D. M. (2024). Fusion algorithms on identifying vacant parking spots using vision-based approach. *Indonesian Journal of Electrical Engineering and Computer Science*, 36(3):1640–1654.

Chen, G.-W., Lin, Y.-H., and İk, T.-U. (2024). Trafficed: Deployment and management system of edge ai cameras. In *NOMS 2024-2024 IEEE Network Operations and Management Symposium*, pages 1–7. IEEE.

Donati, L., Fontanini, T., Tagliaferri, F., and Prati, A. (2020). An energy saving road sweeper using deep vision for garbage detection. *Applied Sciences*, 10(22):8146.

Hannan, M., Al Mamun, M. A., Hussain, A., Basri, H., and Begum, R. A. (2015). A review on technologies and their usage in solid waste monitoring and management systems: Issues and challenges. *Waste Management*, 43:509–523.

Kuo, L.-C. and Lin, H.-Y. (2024). Illegal parking detection based on multi-task driving perception. In *2024 IEEE Intelligent Vehicles Symposium (IV)*, pages 1865–1870. IEEE.

Li, J., Liu, L., Xu, H., Wu, S., and Xue, C. J. (2023). Cross-camera inference on the constrained edge. In *IEEE INFOCOM 2023-IEEE Conference on Computer Communications*, pages 1–10. IEEE.

Liang, Y., Zhang, S., and Wu, J. (2024). Splitstream: Distributed and workload-adaptive video analytics at the edge. *Journal of Network and Computer Applications*, 225:103866.

Matsuda, A., Matsui, T., Matsuda, Y., Suwa, H., and Yasumoto, K. (2021). A method for detecting street parking using dashboard camera videos. *Sensors & Materials*, 33.

Sanchez-Sepulveda, M. V., Navarro, J., Fonseca-Escudero, D., Amo-Filva, D., and Antunez-Anea, F. (2024). Exploiting urban data to address real-world challenges: Enhancing urban mobility for environmental and social well-being. *Cities*, 153:105275.

Tran, D. N.-N., Pham, L. H., Nguyen, H.-H., Tran, T. H.-P., Jeon, H.-J., and Jeon, J. W. (2022). Universal detection-based driving assistance using a mono camera with jetson devices. *IEEE Access*, 10:59400–59412.

Wang, Z., He, X., Zhang, Z., Zhang, Y., Cao, Z., Cheng, W., Wang, W., and Cui, Y. (2022). Edge-assisted real-time video analytics with spatial–temporal redundancy suppression. *IEEE Internet of Things Journal*, 10(7):6324–6335.

Zhao, N., Wang, K., Yang, J., Luan, F., Yuan, L., and Zhang, H. (2024). Cmca-yolo: A study on a real-time object detection model for parking lot surveillance imagery. *Electronics*, 13(8):1557.