SIMON OWEN
JOURNALS | BOOKS | PROCEEDINGS
ACADEMIC PUBLISHERS

Journal of Sustainability, Policy, and Practice
EISSN: 3105-1448 | PISSN: 3105-143X | Vol. 2, No. 1 (2026)

Article

Market-Oriented Perspectives on Dynamic Pricing Decisions under Limited Inventory Conditions

Mason Vance [1,*] and Riley Thorne [2]

[1] Department of Information Technology, George Mason University, Fairfax, USA
[2] College of Engineering, Wayne State University, Detroit, USA
* Correspondence: Mason Vance, Department of Information Technology, George Mason University, Fairfax, USA

Abstract: This review paper investigates the convergence of Unmanned Aerial Vehicle (UAV) technology, semantic segmentation algorithms, and real-time task scheduling on embedded RISC-V platforms. UAVs are increasingly utilized in diverse applications, necessitating efficient onboard processing for tasks such as object detection, environmental mapping, and autonomous navigation. Semantic segmentation, a crucial computer vision technique, enables pixel-level understanding of UAV-captured imagery. However, the computational demands of semantic segmentation algorithms pose a challenge for resource-constrained embedded systems. The RISC-V architecture, an open-source instruction set architecture (ISA), offers a promising solution for developing energy-efficient and customizable hardware platforms for UAVs. This paper provides a comprehensive overview of the current state-of-the-art in UAV semantic segmentation, real-time task scheduling methodologies, and the utilization of RISC-V platforms in this domain. We examine various semantic segmentation algorithms optimized for embedded deployment, focusing on their accuracy, computational complexity, and memory footprint. We also explore different real-time task scheduling techniques employed to manage the execution of semantic segmentation and other critical tasks on UAVs, considering factors such as latency, jitter, and resource utilization. Furthermore, we analyze the advantages and challenges of leveraging RISC-V processors for UAV applications, highlighting their potential for customization, energy efficiency, and security. Finally, we identify key research gaps and future directions in this rapidly evolving field, emphasizing the need for developing novel hardware-software co-design methodologies to enable robust and efficient UAV semantic segmentation on embedded RISC-V platforms. This review contributes to a deeper understanding of the opportunities and challenges in deploying advanced computer vision algorithms on UAVs, facilitating the development of intelligent and autonomous UAV systems.

Keywords: UAV; semantic segmentation; real-time scheduling; RISC-V; embedded systems; computer vision; autonomous navigation

Published: 16 January 2026

1. Introduction

1.1. Motivation and Background

Unmanned Aerial Vehicles (UAVs) are increasingly deployed in diverse applications, including precision agriculture, infrastructure inspection, and disaster response. This widespread adoption necessitates shifting computational tasks from ground stations to onboard embedded systems for enhanced autonomy and reduced latency. A crucial aspect of onboard processing is semantic segmentation, enabling UAVs to understand their environment by classifying each pixel in an image. However, performing complex semantic segmentation in real-time on resource-constrained embedded platforms presents significant challenges. Furthermore, efficient real-time task scheduling is essential to guarantee timely execution of critical tasks, especially when dealing with variable computational loads from algorithms like semantic segmentation, where the processing time t can vary depending on the image complexity c, i.e., $t = f(c)$.

1.2. Problem Statement and Contributions

UAV-based semantic segmentation presents significant challenges when deployed on embedded RISC-V platforms due to limited computational resources and real-time constraints. This paper addresses the problem of efficiently performing semantic segmentation on resource-constrained UAVs using RISC-V processors [1]. We investigate the trade-offs between segmentation accuracy, inference speed, and energy consumption. Our contributions include: (1) a comprehensive analysis of state-of-the-art semantic segmentation models on a RISC-V platform; (2) a novel task scheduling algorithm optimized for real-time semantic segmentation; and (3) an experimental evaluation demonstrating the effectiveness of our approach in achieving high accuracy and low latency for UAV applications, considering the energy budget E, latency L, and accuracy A.

2. Historical Overview of UAV Semantic Segmentation

2.1. Early Approaches to Image Segmentation on UAVs

Early attempts at image segmentation on UAVs primarily relied on traditional computer vision techniques due to the limited computational resources available on early embedded platforms These approaches often involved color-based segmentation, thresholding, and edge detection algorithms [2]. For instance, simple thresholding based on HSV color space values was used to identify vegetation in agricultural applications. Edge detection methods, such as the Sobel operator, were employed to delineate object boundaries. Region growing algorithms, seeded with manually selected points, were also explored to segment images based on pixel similarity. These methods, while computationally inexpensive, often struggled with variations in lighting, shadows, and complex backgrounds, resulting in relatively low segmentation accuracy. Representative early techniques and their characteristics are summarized in Table 1.

Table 1. Early Image Segmentation Techniques for UAVs.

Technique	Description	Limitations
Color-Based Segmentation	Uses color information (e.g., from HSV color space) to identify regions of interest.	Sensitive to lighting changes, shadows, and color variations across different UAV platforms.
Thresholding	Segments images by setting a threshold value. Pixels above/below the threshold are classified into different regions.	Struggles with complex backgrounds and non-uniform illumination. Simple thresholding can only handle bimodal data.
Edge Detection	Identifies object boundaries using operators like the Sobel operator.	Often produces fragmented edges and is sensitive to noise. Requires further processing to close the gaps.
Region Growing	Segments images by iteratively adding neighboring pixels to a seed point, based on similarity criteria.	Requires manual seed point selection and is sensitive to the choice of similarity criteria. Convergence time can be unpredictable.

2.2. Deep Learning-Based Semantic Segmentation

The emergence of deep learning revolutionized semantic segmentation, offering significant improvements over traditional methods. Early deep learning models, such as fully convolutional networks (FCNs), demonstrated the capability to perform pixel-wise classification directly, bypassing the need for hand-crafted features. These FCNs, often based on architectures like VGG or ResNet, were adapted for UAV imagery analysis. Initial applications focused on land cover classification and object detection from aerial perspectives [3]. However, challenges remained in handling the high resolution and viewpoint variations inherent in UAV-captured data, requiring further architectural refinements and specialized training strategies. The key milestones in this evolutionary process are outlined in Table 2.

Table 2. Evolution of Deep Learning Architectures for Semantic Segmentation.

Era	Architecture	Key Features	Application to UAV Imagery	Challenges & Refinements
Early Deep Learning	Fully Convolutional Networks (FCNs)	Pixel-wise classification; End-to-end learning; Adaptation of CNNs like VGG, ResNet	Initial land cover classification and object detection from aerial perspectives	Handling high resolution UAV imagery; Viewpoint variations; Requires architectural refinements and specialized training strategies.
Pre-Deep Learning	Traditional methods (hand-crafted features)	Relied heavily on manually designed features for image analysis.	Used for early land cover mapping and analysis.	Limited by the need for expert knowledge and difficulty in scaling to complex scenes with varying conditions.

2.3. Optimization for Resource-Constrained Environments

Deploying UAV semantic segmentation models on embedded platforms necessitates optimization due to limited resources. Quantization reduces model size and computational complexity by representing weights and activations with lower precision (e.g., INT8 instead of FP32), trading off accuracy for efficiency. Pruning techniques, such as weight or neuron pruning, remove redundant connections, decreasing the number of parameters and operations [4]. Knowledge distillation transfers knowledge from a large, accurate "teacher" model to a smaller, more efficient "student" model, improving the student's performance under resource constraints. These methods enable real-time performance on platforms with limited memory and processing power, making UAV-based applications more practical [5].

3. Semantic Segmentation Algorithms Optimized for UAVs

3.1. Lightweight Convolutional Neural Networks

Lightweight Convolutional Neural Networks (CNNs) are crucial for UAV-based semantic segmentation due to the limited computational resources of embedded platforms. Architectures like MobileNet prioritize efficiency through depthwise separable convolutions, reducing the number of parameters and computational cost compared to standard convolutions. SqueezeNet achieves a small model size by employing fire modules consisting of squeeze and expand layers, significantly decreasing the parameter count while maintaining acceptable accuracy. ShuffleNet further enhances efficiency by utilizing pointwise group convolutions and channel shuffle operations, allowing for information exchange between different channel groups. The performance of these networks on UAV datasets, characterized by aerial imagery and varying perspectives, is often evaluated using metrics like Intersection over Union (IoU) and inference time (t) to determine their suitability for real-time applications [6]. A comparative overview of these lightweight architectures and their performance characteristics is provided in Table 3.

Table 3. Comparison of Lightweight CNN Architectures for UAV Semantic Segmentation.

Architecture	Key Features	Advantages	Disadvantages	Performance Metrics (UAV Datasets)
MobileNet	Depthwise Separable Convolutions	High efficiency (low parameters and computational cost)	Potential accuracy degradation compared to standard CNNs	Evaluated using IoU and inference time t
SqueezeNet	Fire Modules (Squeeze and Expand Layers)	Extremely small model size, reduced parameter count	Can sacrifice some accuracy for size reduction	Evaluated using IoU and inference time t
ShuffleNet	Pointwise Group Convolutions, Channel Shuffle	Enhanced efficiency through information exchange between channel groups	May require careful tuning of group sizes for optimal performance	Evaluated using IoU and inference time t

3.2. Real-Time Semantic Segmentation Techniques

Real-time semantic segmentation on UAVs demands techniques that minimize inference time and maximize frame rates. Several approaches prioritize speed, often involving a trade-off with accuracy. Model compression techniques, such as quantization and pruning, reduce the model size and computational complexity, leading to faster inference. Efficient network architectures, like MobileNetV3 and ShuffleNetV2, are designed with fewer parameters and optimized operations for resource-constrained devices. Furthermore, techniques like knowledge distillation can transfer knowledge from a larger, more accurate model to a smaller, faster one. The acceptable level of accuracy reduction, denoted as ΔA, is application-dependent and must be carefully considered against the gain in frame rate, ΔF [7].

3.3. Domain Adaptation for UAV Imagery

Domain shift poses a significant challenge to deploying semantic segmentation models trained on synthetic or labeled datasets to real-world UAV imagery. This discrepancy arises from differences in image characteristics like resolution, lighting, weather conditions, and sensor noise [8]. Consequently, models trained on one domain often exhibit reduced performance when applied to another. Domain adaptation techniques aim to mitigate this issue by aligning the feature distributions of the source (training) and target (UAV) domains. Common approaches include adversarial training, where a discriminator network distinguishes between source and target features, and feature alignment methods that minimize a distance metric, such as Maximum Mean Discrepancy (MMD), between the domains. These methods enhance the generalization capability of the segmentation model on unseen UAV data, improving its robustness and accuracy [9].

4. Real-Time Task Scheduling on Embedded RISC-V Platforms

4.1. Real-Time Operating Systems (RTOS) for UAVs

Real-Time Operating Systems (RTOS) are crucial for UAVs, enabling deterministic execution of time-critical tasks like flight control and sensor data processing. FreeRTOS, a popular open-source option, offers a small footprint and real-time kernel, suitable for resource-constrained embedded systems. Zephyr, another open-source RTOS, provides a more comprehensive feature set, including support for various communication protocols and security features, making it suitable for more complex UAV applications. Other relevant RTOS options include NuttX, known for its POSIX compliance, and commercial offerings like VxWorks, often chosen for safety-critical applications where certification is required. The selection of an appropriate RTOS depends on factors such as the UAV's computational resources, application complexity, and real-time performance requirements, considering metrics like interrupt latency and task switching overhead ($T_{overhead}$). A structured comparison of these RTOS options is presented in Table 4 [10].

Table 4. Comparison of Real-Time Operating Systems for UAV Applications.

RTOS	Key Features	Suitability for UAVs	Considerations
FreeRTOS	Small footprint, real-time kernel, open-source	Resource-constrained UAVs, basic flight control	Limited features compared to other RTOS options
Zephyr	Comprehensive feature set, communication protocols, security features, open-source	Complex UAV applications, advanced sensor data processing	Larger footprint than FreeRTOS
NuttX	POSIX compliant, open-source	UAVs requiring POSIX compatibility, command-line interfaces	May have a steeper learning curve
VxWorks	Commercial, safety-critical certification	Safety-critical UAV applications, autonomous navigation	Higher cost, licensing restrictions
Interrupt Latency ($T_{latency}$)	Time taken to respond to an interrupt	Real-time performance dependent UAVs	Lower latency RTOS will maintain performance guarantees in time-critical situations
Task Switching Overhead ($T_{overhead}$)	Time taken to switch between different tasks	UAVs with multiple functionalities like simultaneous path planning, control, and environmental perception	Lower $T_{overhead}$ RTOS will ensure real-time performance in multi-tasking situations

4.2. Task Scheduling Algorithms

Task scheduling algorithms are crucial for real-time UAV operation. Rate Monotonic Scheduling (RMS) assigns priorities based on task frequency; tasks with higher frequencies receive higher priorities. Earliest Deadline First (EDF) prioritizes tasks with the nearest deadlines, potentially achieving higher CPU utilization. Priority-based scheduling, a more general approach, allows assigning priorities based on factors beyond frequency or deadline, offering flexibility. For UAVs, RMS is suitable for periodic control tasks with fixed frequencies [11]. EDF can handle aperiodic events like obstacle avoidance, where deadlines are critical. However, EDF's dynamic nature may introduce higher overhead. The choice depends on the specific UAV application and the trade-off between predictability and resource utilization. Priority inversion is a potential issue with all priority-based schemes, requiring mitigation strategies [12].

4.3. Hardware Acceleration for Task Scheduling

Hardware acceleration offers significant potential for enhancing real-time task scheduling on embedded RISC-V platforms. Custom hardware accelerators, designed specifically for scheduling algorithms, can offload computationally intensive tasks from the CPU. This approach reduces the scheduling overhead and improves overall system responsiveness. Field-Programmable Gate Arrays (FPGAs) provide a flexible platform for implementing these accelerators. By configuring the FPGA to execute scheduling functions in parallel, such as priority calculations or deadline comparisons, we can achieve substantial speedups. The performance gain is particularly noticeable when dealing with a large number of tasks (N) and complex scheduling policies, where the computational complexity scales with N or N^2. Furthermore, hardware acceleration can minimize the impact of context switching overhead (T_{cs}), a critical factor in real-time systems [13].

5. Comparison of UAV Semantic Segmentation and Task Scheduling Approaches & Challenges

5.1. Performance Comparison of Semantic Segmentation Models

Different semantic segmentation models exhibit varying performance on UAV datasets. Deep learning models, such as DeepLabv3+ and U-Net, generally achieve higher accuracy, measured by metrics like Intersection over Union (IoU), compared to traditional methods [14]. However, this comes at the cost of increased computational complexity, often quantified by Floating Point Operations Per Second (FLOPS), and larger memory requirements (M). Lightweight models like MobileNetV2-based segmentation networks offer a trade-off, reducing M and FLOPS but potentially sacrificing some accuracy. The choice of model depends on the specific UAV application and the available computational resources [15].

5.2. Challenges in Real-Time Task Scheduling

Ensuring real-time performance for UAV applications presents significant challenges. Task synchronization becomes critical when multiple processes, such as image processing and flight control, must operate in a coordinated manner. Resource contention, particularly for limited memory and processing power on embedded RISC-V platforms, can lead to unpredictable delays. Furthermore, power management is paramount; aggressive power saving strategies can impact task execution times, potentially violating deadlines. Maintaining real-time guarantees while optimizing for energy efficiency requires careful consideration of scheduling algorithms and system-level design. The trade-off between performance and power consumption directly affects the UAV's flight time and operational capabilities, making it a key challenge [16].

5.3. Integration of Semantic Segmentation and Task Scheduling

Integrating semantic segmentation with real-time task scheduling demands careful consideration of computational constraints. One approach involves treating segmentation as a periodic task with a defined deadline. However, variations in processing time due to image complexity can lead to deadline misses. Alternatively, segmentation can be divided into sub-tasks, allowing for preemption and dynamic adjustment of resources based on task criticality. Tradeoffs exist between segmentation accuracy, latency ($T_{latency}$), and resource utilization ($R_{utilization}$). Lowering accuracy, for example, can reduce $T_{latency}$ but impact downstream tasks. Efficient scheduling algorithms are crucial for balancing these competing demands [17].

6. Future Perspectives

6.1. Emerging Trends in UAV Semantic Segmentation

UAV semantic segmentation is poised for significant advancements, driven by innovations in deep learning. Attention mechanisms, allowing networks to focus on salient image regions, will likely become more prevalent, improving segmentation accuracy, especially in complex environments. Transformers, with their ability to model long-range dependencies, offer a promising avenue for enhancing contextual understanding in UAV imagery. Graph Neural Networks (GNNs) can effectively capture relationships between image segments, leading to more coherent and accurate segmentation maps. Furthermore, federated learning presents a compelling solution for training robust semantic segmentation models on UAVs while preserving data privacy. By enabling collaborative learning across multiple UAVs without centralizing data, federated learning can leverage diverse datasets and improve model generalization, addressing the challenges posed by limited on-board data and varying environmental conditions. The parameter t represents the training iteration [18].

6.2. Advancements in RISC-V Architecture

Future RISC-V advancements promise significant benefits for UAV applications. The ongoing development of new extensions, particularly those focused on AI acceleration and real-time processing, will be crucial. We anticipate specialized RISC-V processors tailored for UAVs, incorporating hardware accelerators for tasks like semantic segmentation and path planning [19]. This specialization allows for optimized performance and reduced power consumption, critical for extending flight time. Furthermore, the inherent customizability of RISC-V enables the creation of highly efficient and tailored UAV platforms. By selecting and implementing specific extensions, developers can precisely match the processor's capabilities to the demands of the application, minimizing overhead and maximizing performance per watt. The open nature of RISC-V also fosters innovation, potentially leading to novel architectural solutions for addressing the unique challenges of UAV deployment, such as limited $bandwidth$ and stringent $latency$ requirements [20,21].

7. Conclusion

7.1. Summary of Key Findings

This review highlights the increasing adoption of UAVs for diverse applications, necessitating efficient on-board processing. State-of-the-art UAV semantic segmentation leverages deep learning, often requiring significant computational resources. Real-time task scheduling algorithms, such as EDF and RM, are crucial for ensuring timely execution of critical tasks with deadlines d_i. RISC-V platforms offer a promising open-source alternative for embedded UAV systems, balancing performance and power consumption.

7.2. Concluding Remarks

UAV semantic segmentation on embedded RISC-V platforms presents a promising avenue for real-time applications. However, significant research is still required to optimize models for resource-constrained environments. Future work should focus on developing efficient network architectures, exploring advanced quantization and pruning techniques, and improving task scheduling algorithms to fully leverage the potential of RISC-V based UAVs. The rapidly evolving landscape necessitates continuous innovation in both hardware and software.

References

1. S. Li, K. Liu, and X. Chen, "A context-aware personalized recommendation framework integrating user clustering and BERT-based sentiment analysis," *Journal of Computer, Signal, and System Research*, vol. 2, no. 6, pp. 100-108, 2025.
2. S. Girisha, M. P. MM, U. Verma, and R. M. Pai, "Semantic segmentation of UAV aerial videos using convolutional neural networks," in 2019 IEEE Second International Conference on Artificial Intelligence and Knowledge Engineering (AIKE), 2019, pp. 21-27.
3. B. Jiang, Z. Chen, J. Tan, R. Qu, C. Li, and Y. Li, "A real-time semantic segmentation method based on STDC-CT for recognizing UAV emergency landing zones," Sensors, vol. 23, no. 14, 6514, 2023.
4. X. Zhang, K. Li, Y. Dai, and S. Yi, "Modeling the land cover change in Chesapeake Bay area for precision conservation and green infrastructure planning," Remote Sensing, vol. 16, no. 3, p. 545, 2024. https://doi.org/10.3390/rs16030545
5. Q. Li, H. Yuan, T. Fu, Z. Yu, B. Zheng, and S. Chen, "Multispectral semantic segmentation for UAVs: A benchmark dataset and baseline," IEEE Transactions on Geoscience and Remote Sensing.
6. W. Sun, "Integration of Market-Oriented Development Models and Marketing Strategies in Real Estate," European Journal of Business, Economics & Management, vol. 1, no. 3, pp. 45–52, 2025.
7. S. Yi, J. Li, G. Jiang, X. Liu, and L. Chen, "CCTseg: A cascade composite transformer semantic segmentation network for UAV visual perception," Measurement, vol. 211, 112612, 2023.
8. J. Cheng, C. Deng, Y. Su, Z. An, and Q. Wang, "Methods and datasets on semantic segmentation for Unmanned Aerial Vehicle remote sensing images: A review," ISPRS Journal of Photogrammetry and Remote Sensing, vol. 211, pp. 1-34, 2024.
9. F. Gao, "The role of data analytics in enhancing digital platform user engagement and retention," Journal of Media, Journalism & Communication Studies, vol. 1, no. 1, pp. 10–17, 2025, doi: 10.71222/z27xzp64.

10. Y. Wang, Y. Lyu, Y. Cao, and M. Y. Yang, “Deep learning for semantic segmentation of UAV videos,” in IGARSS 2019-2019 IEEE International Geoscience and Remote Sensing Symposium, 2019, pp. 2459-2462.
11. M. K. Masouleh and R. Shah-Hosseini, “Development and evaluation of a deep learning model for real-time ground vehicle semantic segmentation from UAV-based thermal infrared imagery,” ISPRS Journal of Photogrammetry and Remote Sensing, vol. 155, pp. 172-186, 2019.
12. C. L. Cheong, “Research on AI Security Strategies and Practical Approaches for Risk Management”, J. Comput. Signal Syst. Res., vol. 2, no. 7, pp. 98–115, Dec. 2025, doi: 10.71222/17gqja14.
13. S. Liu, J. Cheng, L. Liang, H. Bai, and W. Dang, “Light-weight semantic segmentation network for UAV remote sensing images,” IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing, vol. 14, pp. 8287-8296, 2021.
14. G. Ying, “Machine learning and cloud-enhanced real-time distributed systems for intelligent urban services,” Journal of Science, Innovation & Social Impact, vol. 1, no. 1, pp. 189–200, 2025.
15. J. Deng, Z. Zhong, H. Huang, Y. Lan, Y. Han, and Y. Zhang, “Lightweight semantic segmentation network for real-time weed mapping using unmanned aerial vehicles,” Applied Sciences, vol. 10, no. 20, 7132, 2020.
16. S. Yuan, “Data Flow Mechanisms and Model Applications in Intelligent Business Operation Platforms”, Financial Economics Insights, vol. 2, no. 1, pp. 144–151, 2025, doi: 10.70088/m66tbm53.
17. A. S. Chakravarthy, S. Sinha, P. Narang, M. Mandal, V. Chamola, and F. R. Yu, “DroneSegNet: Robust aerial semantic segmentation for UAV-based IoT applications,” IEEE Transactions on Vehicular Technology, vol. 71, no. 4, pp. 4277-4286, 2022.
18. S. A. Ahmed, H. Desa, H. K. Easa, A. S. T. Hussain, T. A. Taha, S. Q. Salih, et al., “Advancements in UAV image semantic segmentation: A comprehensive literature review,” Multidisciplinary Reviews, vol. 7, no. 6, 2024118-2024118, 2024.
19. L. U. Xudong, W. U. Yiquan, and C. H. E. N. Jinlin, “Research progress on deep learning methods for object detection and semantic segmentation in UAV aerial images,” Acta Aeronautica et Astronautica Sinica, vol. 45, no. 6, 2024.
20. R. Luo, X. Chen, and Z. Ding, "SeqUDA-Rec: Sequential user behavior enhanced recommendation via global unsupervised data augmentation for personalized content marketing," *arXiv preprint* arXiv:2509.17361, 2025.
21. Y. Chen, H. Du, and Y. Zhou, “Lightweight network-based semantic segmentation for UAVs and its RISC-V implementation,” Journal of Technology Innovation and Engineering, vol. 1, no. 2, 2025.