FloodNav: Liquid Neural Networks with Reinforcement Learning for Precision Drone Flight in Disaster Zones

Yash Potdar[1], Sahil Pawar[2], Ganesh Wani[3], Akanksha Singh[4]

[1,2,3,4]Department of AI and Data Science, AISSMS IOIT

Abstract

Unmanned Aerial Vehicles (UAVs) are increasingly being deployed in applications such as surveillance, transportation, disaster response, and smart city infrastructure. Their autonomy relies heavily on two key capabilities: reliable navigation through dynamic and unpredictable environments, and real-time object detection using resource-constrained on- board hardware.

This review paper examines advancements in global and local path planning algorithms, with particular emphasis on A* and Timed Elastic Band (TEB), which together enable both efficient global trajectory generation and adaptive local obstacle avoidance [1]. In parallel, the paper explores embedded artificial intelligence approaches that leverage lightweight YOLO architectures on edge platforms such as Raspberry Pi, emphasizing optimization techniques including quantization, pruning, and hardware acceleration [2, 3].

Comparative analysis reveals that hybrid planning strategies significantly improve trajectory efficiency, while quantized and hardware-accelerated YOLO models achieve real- time inference with reduced power consumption. However, challenges persist in thermal management, dataset generalization, and achieving consistent performance across heterogeneous hardware.

By synthesizing the latest research trends, this review highlights the convergence of robust navigation frameworks with optimized lightweight deep learning models as the foundation for next-generation UAVs. These advancements mark a decisive step toward fully autonomous systems that are efficient, adaptive, and reliable in real-world missions. [4, 5]

Keywords: UAV, Path Planning, YOLO, Edge AI, Raspberry Pi, Embedded Systems

1 Introduction

The rapid proliferation of unmanned aerial vehicles (UAVs) is revolutionizing applications across transportation, surveillance, emergency response, and smart city infrastructure [1, 8, 9]. Modern UAVs rely on advanced autonomy, requiring robust solutions for real-time navigation and environmental perception. Central to their operation are two intersecting challenges: planning collision-free trajectories in complex, unpredictable environments, [4] and accurately recognizing critical objects on-the-fly using resource-constrained hardware.

The rapid proliferation of unmanned aerial vehicles (UAVs) is revolutionizing applications across transportation, surveillance, emergency response, and smart city infrastructure [1]. Mod- ern UAVs rely on advanced autonomy, requiring robust solutions for real-time navigation and environmental perception. Central to their operation are two intersecting challenges: planning collision-free trajectories in complex, unpredictable environments, and accurately recognizing critical objects on-the-fly using resource-constrained hardware.

Effective path planning allows drones to safely maneuver through dynamic spaces, avoiding both static and moving obstacles. Classical algorithms, like A-star, deliver reliable global routes based on prior environmental knowledge, while contemporary techniques such as the Timed Elastic Band (TEB) algorithm offer flexible, adaptive local planning, enabling real-time obstacle avoidance and trajectory refinement [1]. The integration of these approaches within open-source frameworks like ROS, coupled with sensor modalities such as LiDAR, empowers autonomous UAVs to operate effectively in both indoor and outdoor environments, often where traditional satellite-based navigation like GPS is unavailable.

Simultaneously, the emergence of embedded artificial intelligence has enabled UAVs to execute complex visual recognition tasks directly on onboard platforms like Raspberry Pi [2, 3]. Lightweight, quantized deep learning models—most notably the YOLO family—support high-speed, power-efficient inference for aerial object detection, crucial in safety-critical and time-sensitive missions. Deploying such models on edge devices introduces new constraints in computational capacity, thermal management, and energy consumption, while demanding robust detection accuracy in diverse real-world conditions.

This paper presents a comprehensive review of topological path planning algorithms and edge AI deployment strategies for autonomous drones. By critically analyzing advances in trajectory optimization and real-time embedded object detection, the study aims to synthesize current best practices and identify future directions towards fully autonomous, intelligent UAV systems that combine agility, reliability, and efficiency.

2 Literature Review

Research on UAV autonomy has evolved along two major dimensions: (i) **navigation and path planning**, focusing on obstacle avoidance and efficient trajectory generation, and (ii) **embedded object detection**, leveraging lightweight deep learning models optimized for real-time inference on resource-constrained platforms. This section reviews key works representing both domains.

2.1 Navigation and Path Planning for UAVs

Path planning remains one of the most critical challenges for UAV autonomy, requiring robust algorithms to generate collision-free trajectories in dynamic and uncertain environments. Classical search-based algorithms such as A* have been widely adopted due to their simplicity, deterministic nature, and ability to find globally optimal solutions in grid-based maps. However, their reliance on static environmental models limits adaptability in real-world scenarios [1].

To address dynamic conditions, local planners such as the *Timed Elastic Band* (TEB) algorithm have been integrated into UAV navigation frameworks. TEB optimizes trajectories by considering both geometric and temporal constraints, enabling real-time adjustments to avoid moving obstacles while maintaining smooth flight dynamics [1]. Recent works demonstrate the integration of A* as a global planner and TEB as a local planner within the Robot Operating System (ROS), supported by sensor modalities such as LiDAR and RGB-D cameras. Simulation studies in Gazebo environments show that TEB outperforms A* by approximately 10.6% in trajectory efficiency within a 30m range, highlighting the effectiveness of hybrid global-local planning strategies for UAVs. [6, 7]

2.2 Embedded Object Detection using YOLOv4-Tiny

Parallel to navigation, embedded artificial intelligence has transformed UAV perception, particularly in safety-critical applications such as emergency response. YOLO (You Only Look Once) models are well-known for their real-time detection capabilities. However, their deployment on resource-limited platforms like Raspberry Pi requires optimization to reduce computational overhead.

Boddu and Mukherjee [2] proposed an efficient edge deployment of a quantized YOLOv4- Tiny model for aerial emergency object detection. The model was optimized via TensorFlow Lite post-training quantization, reducing weights from FP32 to INT8 precision. The quantized model achieved an inference time of 183 ms per image on Raspberry Pi 5, a 36% improvement compared to its FP32 counterpart. Power consumption decreased by 43.9% (from 7.13 W to 4.0 W), while maintaining robust detection accuracy across emergency classes such as ambulance, fire engine, and car crash. [8] These results underscore the viability of deploying quantized lightweight detectors for embedded UAV applications where battery efficiency is paramount.

2.3 UAV Detection with YOLOv8 on Raspberry Pi

More recently, research has advanced towards state-of-the-art YOLOv8 architectures, which offer enhanced accuracy and flexibility. Kryvenchuk et al. [3] investigated the feasibility of deploying YOLOv8 models on Raspberry Pi 5, integrated with a Pi Camera Module 3, for standalone UAV detection. Multiple configurations were evaluated, including YOLOv8n, YOLOv8 with Selective Knowledge Distillation (SKD), and YOLOv8 accelerated with the HailoAI hard- ware module.

Empirical findings show that the baseline YOLOv8 achieved the highest mean Average Precision (mAP) of 87.2% but was limited to 7.4 FPS, insufficient for fast UAV maneuvers. YOLOv8n, with reduced complexity, improved speed to 13.2 FPS while sacrificing some ac- curacy (82.6% mAP). The SKD-enhanced YOLOv8 balanced accuracy and efficiency (85.3% mAP at 10.1 FPS). The HailoAI-assisted version demonstrated the best trade-off, achieving 35 FPS with a mAP of 86.9%, significantly improving real-time performance while managing power consumption effectively.

2.4 Summary of Literature

The reviewed studies demonstrate clear trends in UAV autonomy. Navigation-focused research prioritizes hybrid planning frameworks (A* + TEB) that balance global optimality with local adaptability [1]. Detection-focused works emphasize lightweight, quantized, and hardware- accelerated YOLO models to ensure real-time inference within the power and thermal limits of embedded devices [2, 3]. Together, these efforts highlight the convergence of robust trajectory planning and efficient onboard perception as the foundation for next-generation intelligent UAV systems.

3 Methodology

This review employs a structured methodology to systematically identify, select, and synthesize relevant literature on UAV navigation and real-time detection using embedded artificial intelligence (AI). The process is divided into four stages: literature search, selection criteria, data extraction, and comparative synthesis.

3.1 Literature Search Strategy

A comprehensive search was conducted across academic databases, including *IEEE Xplore, arXiv, Springer, and CEUR Workshop Proceedings*, as well as robotics and AI conference proceedings. Boolean keyword combinations such as *"UAV navigation," "autonomous drones," "ROS path planning," "YOLO," "Raspberry Pi," "edge AI," "quantization,"* and *"embedded object detection"* were employed. This ensured coverage of both navigation-focused research and object detection studies optimized for resource-constrained platforms [1, 2, 3].

3.2 Selection Criteria

From the initial corpus of literature, studies were screened using defined inclusion and exclusion criteria. The **inclusion criteria** consisted of:

- Research addressing UAV navigation using algorithms such as A* and Timed Elastic Band (TEB) [1].
- Deployment of lightweight YOLO models (e.g., YOLOv4-Tiny, YOLOv8) on embedded platforms such as Raspberry Pi 5 [2, 3].
- Studies reporting evaluation metrics including accuracy, inference time, power efficiency, and system latency.

The **exclusion criteria** included:

- Works limited to theoretical frameworks or simulations requiring high-end GPUs, with- out embedded implementation.
- Non-peer-reviewed sources or studies lacking quantitative evaluation.

Following this screening process, three representative studies were shortlisted for in-depth review:

1. Online planning of topologically distinctive autonomous drone trajectories [1].
2. Efficient edge deployment of quantized YOLOv4-Tiny for aerial emergency detection [2].
3. UAV detection with YOLOv8 on Raspberry Pi 5 [3].

3.3 Data Extraction and Categorization

Each study was analyzed to extract key information, categorized under the following dimensions:

- **Algorithms and Models:** Path planning approaches (A*, TEB) and YOLO variants (YOLOv4-Tiny, YOLOv8).
- **Hardware Platforms:** Raspberry Pi 5, Pixhawk flight controllers, LiDAR, and HailoAI accelerators.
- **Optimization Techniques:** Quantization (FP32 to INT8), pruning, and Selective Knowledge Distillation (SKD).
- **Evaluation Metrics:** Mean Average Precision (mAP), Precision, Recall, F1-score, Frames per Second (FPS), inference latency, and power consumption.

3.4 Comparative Synthesis

The extracted data was then systematically compared to highlight trade-offs across approaches. Navigation-focused research emphasized the integration of global and local planning for adapt- ability in dynamic environments [1], while detection-focused studies optimized YOLO architectures for real-time inference on embedded devices [2, 3]. Special attention was given to the impact of quantization, pruning, and hardware acceleration in reducing computational overhead while maintaining detection accuracy.

This structured approach allows the review to synthesize advancements in UAV navigation and detection, while identifying gaps such as thermal profiling of embedded platforms, sensor fusion for robust UAV detection, and cross-platform benchmarking of optimized AI models.

4 Comparative Analysis and Discussion

A critical comparison of the reviewed studies highlights the trade-offs between navigation algorithms and embedded object detection models. While navigation-focused approaches prioritize trajectory efficiency and adaptability, detection-focused studies optimize inference speed and energy efficiency on constrained platforms such as Raspberry Pi.

Table 1 summarizes the major contributions, evaluation metrics, and observed trade-offs across the three-representative works.

From Table 1, several insights emerge:

- **Navigation trade-offs:** A* provides globally optimal paths but lacks adaptability, while TEB enables real-time obstacle avoidance, improving trajectory efficiency in dynamic environments.
- **Detection efficiency:** Quantized YOLOv4-Tiny achieved substantial reductions in power consumption and inference time, proving effective for battery-powered UAV missions.
- **Real-time feasibility:** YOLOv8 with HailoAI acceleration demonstrated the highest frame rates (35 FPS) with near-state-of-the-art accuracy, showing that hardware accelerators can overcome embedded platform constraints.

Overall, the reviewed literature suggests that combining hybrid planning frameworks with lightweight, optimized YOLO detectors offers a promising pathway towards fully autonomous UAV systems that are both agile in navigation and reliable in perception.

5 FloodNav: Project Differentiation and Expected Results

Building upon the reviewed advancements in UAV navigation and embedded perception, our project **FloodNav — Liquid Neural Networks with Reinforcement Learning for Precision**

Table 1: Comparative analysis of UAV navigation and embedded detection approaches.

Study	Approach / Model	Accuracy / mAP	FPS	Power	Key Contribution
Vairagi et al. (2024)	A* (global) + TEB (local) planners	--	Real-time re-planning	--	TEB improved trajectory efficiency by ~10.6% over A* in 30m test environments
Boddu & Mukherjee (2024)	YOLOv4-Tiny (INT8 quantized) on Raspberry Pi 5	High detection accuracy maintained	~5.5	4.0 W (43.9% lower than FP32)	Quantization reduced model size (22.5 MB → 6.4 MB), inference time (36% faster), and power consumption
Kryvenchuk et	YOLOv8	YOLOv8: 87.2% •	7.4 / 13.2 /	--	HailoAI

al. (2025)	variants on Raspberry Pi 5	YOLOv8n: 82.6% • YOLOv8+SKD: 85.3% • YOLOv8+HailoAI: 86.9%	10.1 / 35		acceleration offered best trade-off: high mAP (86.9%) with real-time speed (35 FPS)

Drone Flight in Disaster Zones introduces a novel and mission-critical application of autonomous aerial systems in flood response scenarios. While existing research has laid strong foundations in hybrid path planning (A* + TEB) and edge-optimized object detection (e.g., YOLOv4-Tiny, YOLOv8), these methods often target generic navigation or static perception problems. In contrast, FloodNav addresses the highly dynamic, uncertain, and time-sensitive context of **large-scale flood disasters**, where rapid situational awareness and precise aerial operations can save lives.

5.1 Motivation

Flood disasters frequently disrupt ground transportation, communication networks, and power grids, making rapid aerial mapping and monitoring crucial for effective disaster response. [8, 9] Traditional UAV solutions relying on static path planning or generic object detection often fail to adapt to rapidly evolving flood boundaries, submerged landscapes, or unexpected obstacles such as floating debris. Additionally, flight missions during disasters are constrained by limited battery power, unpredictable weather, and the need for accurate, real-time geospatial intelligence.

FloodNav is designed to overcome these challenges by integrating:

- **Adaptive visual intelligence** through liquid neural networks to perceive changing flood extents robustly.
- **Reinforcement learning-based path optimization** to dynamically adjust drone trajectories for efficient coverage of critical zones.
- **Geospatial prioritization** to focus UAV resources on high-population and high-risk areas.
- **Real-time decision dashboards** to translate UAV insights into actionable strategies for emergency teams.

5.2 Methodological Differentiation from Existing Work

Compared to the surveyed literature, FloodNav introduces several methodological innovations across perception, navigation, and operational deployment:

5.2.1 1. Perception: Liquid Neural Networks for Flood Detection

Previous research on onboard perception has primarily focused on deploying lightweight YOLO architectures on embedded platforms, often optimized through quantization and pruning [2, 3]. While these methods achieve high inference speeds and energy efficiency, they are typically trained on static datasets for generic object categories (e.g., vehicles, people), and may under- perform in unfamiliar or evolving environments.

FloodNav replaces traditional CNNs with **liquid neural networks (LNNs)**, a new class of dynamic networks that adapt their internal dynamics to incoming data streams in real time. This allows the perception module to:

- Continuously adjust to changing flood conditions, such as rising water levels or varying illumination.
- Operate effectively with fewer parameters, improving energy efficiency while maintaining adaptability.
- Provide temporally aware segmentation and flood boundary detection without frequent model retraining.

This adaptive capability is critical in disaster zones, where the visual environment can change significantly between flights.

5.2.2 2. Navigation: Reinforcement Learning for Path Optimization

Existing navigation frameworks integrate global A* and local TEB planners to balance global optimality with local adaptability [1]. These methods perform well in structured environments but can struggle with complex, evolving terrains where obstacles emerge unpredictably.

FloodNav employs **reinforcement learning (RL)** to train UAV agents to autonomously learn optimized flight strategies. Through reward functions designed around coverage efficiency, energy consumption, and priority zone focus, the RL agent learns to:

- Minimize flight time and battery usage while maximizing the coverage of critical flood- affected areas.
- Adapt paths dynamically to avoid newly detected hazards and flooded regions.
- Continuously improve performance through simulation-to-real transfer using realistic flood environment models. [6, 7]

5.2.3 3. Operational Context: Geospatial Intelligence and Disaster Dashboard

Unlike prior works that focus on either navigation or detection in isolation, FloodNav integrates both into a **real-time operational dashboard**. This interface provides emergency teams with:

- Live drone telemetry, including flight paths, energy usage, and current location.
- Dynamic flood maps generated by onboard LNN perception.
- Geofencing alerts for newly identified hazardous or restricted areas.
- Suggested optimized flight missions based on reinforcement learning outputs.

This integration transforms UAV data into actionable disaster intelligence, enabling faster response coordination.

5.3 Expected Results

Based on the literature review and the innovative methodology, the following outcomes are expected from FloodNav:

- **Improved Flight Efficiency:** RL-driven flight planning is projected to reduce mission time and energy consumption compared to classical A*+TEB approaches, especially in dynamically changing terrains. [6] This improvement could extend operational endurance by 15–25% in simulated flood scenarios.
- **Enhanced Perception Robustness:** Liquid neural networks are expected to outperform static YOLO-based models in flood boundary detection, particularly under conditions of lighting variability, occlusion, and rapidly changing landscapes.

- **Real-Time Situational Awareness:** The integration of onboard intelligence with a centralized dashboard will enable live monitoring and mission adjustment, improving response times for emergency teams.
- **Scalability and Adaptability:** The modular design of FloodNav allows deployment on diverse UAV platforms, supporting both edge computing (e.g., Raspberry Pi, Jetson) and cloud offloading when available, ensuring flexible operation in different disaster zones.

5.4 Potential Impact

FloodNav has the potential to significantly enhance disaster response capabilities by:

- Reducing the time required to map flood extents and identify critical zones from hours to minutes.
- Prioritizing areas with high population density and infrastructure value, leading to more effective allocation of rescue resources.
- Enabling multiple UAVs to operate collaboratively, scaling coverage across large flood- affected regions.
- Providing a replicable and extensible framework for other disaster contexts, such as wild- fires or landslides.

In summary, FloodNav represents a substantial advancement over existing UAV navigation and detection frameworks by integrating adaptive neural models, learning-based flight strategies, and real-time operational intelligence in a disaster response setting.

6 Challenges and Future Directions

Despite significant progress in UAV navigation and embedded object detection, several challenges remain that hinder the deployment of fully autonomous systems in real-world environments.

6.1 Challenges

- **Computational and thermal constraints:** Raspberry Pi and similar edge devices face heating issues during prolonged inference, which can lead to throttling and degraded performance.
- **Energy limitations:** UAVs operate under strict battery budgets, and higher computation demands reduce flight duration. Balancing accuracy with power efficiency remains a critical challenge.
- **Dataset diversity:** Many detection models are trained on limited datasets that may not capture real-world conditions such as varying weather, lighting, or occlusions, leading to reduced robustness.
- **Navigation in dynamic environments:** While hybrid planners like A*+TEB improve adaptability, they still struggle in highly dynamic and cluttered scenarios where multiple moving obstacles interact.
- **Hardware dependency:** Performance improvements often rely on specialized accelerators (e.g., HailoAI), which may limit portability and increase system cost.

6.2 Future Directions

- **Sensor fusion:** Integrating LiDAR, RGB-D cameras, and inertial sensors can enhance perception accuracy and provide redundancy in GPS-denied environments.
- **Energy-aware navigation:** Developing trajectory planners that incorporate energy consumption models will help maximize UAV flight endurance.
- **Cross-platform benchmarking:** Standardized benchmarks across devices such as Rasp- berry Pi,

NVIDIA Jetson, and Google Coral are needed to evaluate trade-offs in performance and efficiency.

- **Hybrid optimization methods:** Combining quantization, pruning, and knowledge distillation may unlock higher compression rates while maintaining near-original accuracy.
- **Collaborative UAV systems:** Future research may explore multi-UAV networks that coordinate navigation and detection tasks, enhancing scalability for large-area coverage.

7 Conclusion

This review has synthesized recent advances in UAV navigation and embedded object detection, focusing on two intersecting challenges: generating efficient trajectories in dynamic environments and executing accurate real-time perception on resource-constrained platforms.

For navigation, the integration of global planners like A* with local adaptive strategies such as the Timed Elastic Band (TEB) algorithm has been shown to improve trajectory efficiency and obstacle avoidance performance [1]. For perception, the deployment of lightweight YOLO architectures on embedded platforms such as Raspberry Pi demonstrates that quantization, pruning, and hardware acceleration can significantly reduce computational overhead without severely impacting detection accuracy [2, 3].

The comparative analysis highlights that while classical algorithms ensure reliability, mod- ern approaches prioritize adaptability and efficiency. Similarly, YOLOv8 combined with hardware accelerators such as HailoAI has emerged as a promising solution for real-time UAV detection, achieving near state-of-the-art accuracy with practical inference speeds.

Despite these advancements, challenges remain in thermal management, energy constraints, dataset diversity, and hardware portability. Addressing these issues will require a shift toward holistic strategies that combine sensor fusion, energy-aware planning, hybrid optimization, and collaborative multi-UAV systems.

In conclusion, the convergence of robust navigation frameworks with optimized lightweight AI models represents a decisive step toward building the next generation of intelligent, autonomous UAVs that are agile, efficient, and reliable in real-world missions.

References

[1] R. D. Vairagi, V. B. Semwal, M. Vishwakarma, “Online Planning of Topologically Distinctive Autonomous Drone Trajectories,” IEEE International Students’ Conference on Electrical, Electronics and Computer Science, 2024.

[2] S. Boddu, A. Mukherjee, “Efficient Edge Deployment of Quantized YOLOv4-Tiny for Aerial Emergency Object Detection on Raspberry Pi 5,” 2024.

[3] Y. Kryvenchuk, R. Stupnytskyi, et al., “UAV Detection with YOLO on a Standalone Rasp- berry Pi 5 System,” 2nd International Conference on Smart Automation & Robotics for Future Industry, 2025.

[4] R. Hasani, D. Dey, and D. Rus, “Liquid Neural Networks for Adaptive Control of Autonomous Robots,” Nature Machine Intelligence, vol. 5, pp. 1180–1192, 2023.

[5] Y. Wang, T. Liu, et al., “Adaptive Aerial Navigation using Liquid Time-Constant Net- works,” IEEE Transactions on Neural Networks and Learning Systems, 2024.

[6] K. Singh and P. Sharma, “Deep Reinforcement Learning-based UAV Path Planning for Dynamic Disaster Environments,” IEEE Access, vol. 12, pp. 75849–75861, 2024.

[7] L. Zhang, Q. Chen, et al., “Multi-Agent Reinforcement Learning for Cooperative Aerial Disaster

Mapping," Robotics and Autonomous Systems, vol. 180, 104763, 2025.

[8] M. Alam, S. K. Roy, et al., "Autonomous UAV Systems for Flood Monitoring and Early Warning: A Deep Learning Perspective," Journal of Field Robotics, 2023.

[9] R. Patel, A. Nair, et al., "Aerial Flood Assessment Using Edge-AI Drones and Cloud Collaboration," IEEE Transactions on Geoscience and Remote Sensing, 2024.