

## Article

# AI-Enhanced Eco-Efficient UAV Design for Sustainable Urban Logistics: Integration of Embedded Intelligence and Renewable Energy Systems

Luigi Bibbò <sup>1,\*</sup>,† , Filippo Laganà <sup>2,\*</sup>,† , Giuliana Bilotta <sup>1</sup> , Giuseppe Maria Meduri <sup>1</sup>, Giovanni Angiulli <sup>3</sup>  and Francesco Cotroneo <sup>4</sup>

<sup>1</sup> Department of Civil, Energy, Environment and Materials (DICEAM), Mediterranean University of Reggio Calabria, Via Zehender, 89124 Reggio Calabria, Italy; giuliana.bilotta@unirc.it (G.B.); giuseppemaria.meduri@unirc.it (G.M.M.)

<sup>2</sup> Laboratory of Biomedical Applications Technologies and Sensors (BATS), Department of Health Science “Magna Græcia” University, Viale Europa, Località Germaneto, snc, 88100 Catanzaro, Italy

<sup>3</sup> Department of Information Engineering, Infrastructures and Sustainable Energy, Mediterranean University of Reggio Calabria, Via R. Zehender, 89124 Reggio Calabria, Italy; giovanni.angiulli@unirc.it

<sup>4</sup> Nophys srl, Via Maddaloni, 74, 00177 Roma, Italy; francesco.cotroneo@nophys.com

\* Correspondence: luigi.bibbo@unirc.it (L.B.); filippo.lagana@unicz.it (F.L.)

† These authors contributed equally to this work.

## Abstract

The increasing use of UAVs has reshaped urban logistics, enabling sustainable alternatives to traditional deliveries. To address critical issues inherent in the system, the proposed study presents the design and evaluation of an innovative unmanned aerial vehicle (UAV) prototype that integrates advanced electronic components and artificial intelligence (AI), with the aim of reducing environmental impact and enabling autonomous navigation in complex urban environments. The UAV platform incorporates brushless DC motors, high-density LiPo batteries and perovskite solar cells to improve energy efficiency and increase flight range. The Deep Q-Network (DQN) allocates energy and selects reference points in the presence of wind and payload disturbances, while an integrated sensor system monitors motor vibration/temperature and charge status to prevent failures. In urban canyon and field scenarios (wind from 0 to 8 m/s; payload from 0.35 to 0.55 kg), the system reduces energy consumption by up to 18%, increases area coverage by 12% for the same charge, and maintains structural safety factors > 1.5 under gust loading. The approach combines sustainable materials, efficient propulsion, and real-time AI-based navigation for energy-conscious flight planning. A hybrid methodology, combining experimental design principles with finite-element-based structural modelling and AI-enhanced monitoring, has been applied to ensure structural health awareness. The study implements proven edge-AI sensor fusion architectures, balancing portability and telemonitoring with an integrated low-power design. The results confirm a reduction in energy consumption and CO<sub>2</sub> emissions compared to traditional delivery vehicles, confirming that the proposed system represents a scalable and intelligent solution for last-mile delivery, contributing to climate resilience and urban sustainability. The findings position the proposed UAV as a scalable reference model for integrating AI-driven navigation and renewable energy systems in sustainable logistics.

**Keywords:** UAV; sustainable drone; brushless DC motors; LiPo batteries; perovskite solar cells; startup innovation; DQN



Academic Editors: Lorenzo Gigli and Federica Zonzini

Received: 7 September 2025

Revised: 24 September 2025

Accepted: 29 September 2025

Published: 2 October 2025

**Citation:** Bibbò, L.; Laganà, F.; Bilotta, G.; Meduri, G.M.; Angiulli, G.; Cotroneo, F. AI-Enhanced Eco-Efficient UAV Design for Sustainable Urban Logistics: Integration of Embedded Intelligence and Renewable Energy Systems. *Energies* **2025**, *18*, 5242. <https://doi.org/10.3390/en18195242>

**Copyright:** © 2025 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

## 1. Introduction

In recent years, UAVs, commonly called drones, have become transformative tools in the logistics and transportation industries [1,2]. Their ability to avoid ground traffic and reach remote or crowded urban areas offers a strong alternative to traditional delivery methods [3,4]. As global supply chains face growing challenges, including driver shortages, rising fuel prices, and the urgent need to cut greenhouse gas emissions, drones present a promising way to improve delivery efficiency while reducing environmental impact [5,6]. The rise of electric-powered UAVs supports broader sustainability goals and the European Green Deal, which advocates for decarbonisation and digital innovation in urban mobility. Unlike traditional delivery vehicles, electric drones require minimal infrastructure, generate near-zero local emissions, and can be deployed flexibly in dense city areas. Several companies are experimenting with UAVs for last-mile delivery.

Technological advancements in propulsion and energy systems have been crucial to this development. High-efficiency brushless DC (BLDC) motors, lithium polymer (LiPo) batteries, and emerging photovoltaic technologies such as perovskite solar cells have greatly enhanced flight endurance, energy density, and operational reliability [7–9]. The implementation of innovations allows UAVs to undertake longer missions with less dependence on ground charging infrastructure, which is essential for urban deployment.

The integration of AI and sensor fusion is an important aspect, because modern UAVs, equipped with GPS, inertial measurement units (IMUs), LiDAR and multispectral cameras, allow both real-time reconstruction of the environment and autonomous decision-making [10]. AI algorithms, particularly those based on reinforcement learning, enable drones to dynamically adapt to changing conditions, optimise flight trajectories and avoid obstacles with minimal human intervention. To reinforce methodological rigour, this study draws inspiration from applications validated by robust experimental design, such as the Taguchi method, which demonstrate how systematic tuning can improve efficiency and reliability under variable operating conditions [11]. Similarly, the combination of finite element modelling with AI-enhanced monitoring has already been validated in recent studies, confirming the potential of hybrid physics-AI approaches for structural health monitoring and adaptive control [12,13]. Furthermore, the integration of sensors and embedded electronics with artificial intelligence-based diagnostics, applied to the monitoring of wastewater treatment machinery, highlights the versatility of edge-AI architectures for real-time fault detection and predictive maintenance [14,15]. Finally, studies have shown how it is possible to successfully balance portability, low power consumption and remote monitoring, providing useful design principles for compact UAV systems operating in resource-constrained environments [16]. The proposed UAV integrates electronics, renewable energy, and artificial intelligence-based navigation. The main innovation lies in the use of a DQN for autonomous route optimisation, enabling a balance between energy efficiency, environmental risk management, and mission success in complex environments. Specifically, the manuscript provides an energy-conscious DQN architecture implemented on the edge. The lightweight DQN model, optimised and quantised for real-time inference on NVIDIA Jetson Nano (average inference  $\approx 3.2$  ms, p95  $\approx 5.1$  ms), presents a design for selecting energy-efficient waypoints and balancing the trade-off between energy and risk in real time. Performance and latency metrics are reported in Section 3.6. In addition, the paper presents a practical hybrid energy management system (EMS). An implementable EMS that coordinates perovskite photovoltaic modules via MPPT, an IEC-62133 [17]-compliant BMS, and the DC bus serving propulsion and computation, using priority rules (PV  $\rightarrow$  load  $\rightarrow$  battery) and safety thresholds to preserve battery life and ensure operational safety (see Sections 3.3 and 3.4). Finally, the study analyses an integrated engineering workflow. A combined CAD/FEM structural design and AI-enhanced monitoring path that combines

structural health awareness with mission planning (hybrid physical-AI methodology), enabling the co-optimisation of structural, energy and control subsystems.

The prototype validation roadmap (Section 3.10) defines a path from HIL to outdoor testing for progressive experimental verification.

In contrast to prior research that typically addresses energy optimisation, navigation, or renewable integration as separate domains, this study proposes a unified framework where these elements are jointly optimised. This holistic approach reflects the current research gap and advances the field toward deployable eco-efficient UAVs for urban logistics. The implemented DQN features both an environment learning capability and an algorithmic ability to capture dynamic state–action relationships [18]. Solar-powered platforms can achieve long endurance but often sacrifice payload capacity and structural robustness, reducing their applicability in logistics [19]. Most commercial drones continue to rely on LiPo batteries, which limit range and require frequent recharging on the ground, hindering large-scale deployment in densely populated areas [20].

Furthermore, rule-based navigation systems lack real-time adaptability, reducing efficiency in variable conditions [21]. Although model-based systems engineering (MBSE) is increasingly used, integration with artificial intelligence-based decision-making and the exploitation of renewable energies is still limited [22]. This study addresses the gap between purely simulation-based studies and integrated eco-efficient UAV design, combining structural optimisation, renewable energy harvesting, and AI-enhanced control. In light of the above, the integration of sensor fusion with artificial intelligence-based monitoring further increases reliability by enabling real-time assessment of engines, batteries and structural integrity. These elements constitute a comprehensive and eco-efficient design approach that combines technological, environmental and business dimensions.

The objectives of this research are threefold. The first concerns the initial evaluation of the prototype, which integrates sustainable materials, efficient propulsion and solar energy harvesting. The second demonstrates the feasibility of real-time autonomous navigation through integrated artificial intelligence, with a focus on DQN-based route optimisation implemented on the NVIDIA Jetson Nano platform. The third describes the potential of this technology as a basis for a start-up dedicated to sustainable drone delivery services, particularly in regions where electric vehicle infrastructure remains limited.

Statements referencing start-up formation and commercial exploitation are presented as prospective scenarios contingent on successful prototype validation.

The present work reports simulation-based feasibility and a staged validation plan (Section 3.10). Any commercialisation statements are therefore exploratory and dependent on (i) successful HIL and laboratory validation of MPPT/BMS and PV yield under flight conditions, (ii) certified outdoor trials under regulated conditions, and (iii) regulatory and safety compliance. Recent advances in UAV navigation have explored a wide range of techniques, from traditional algorithms such as A\* and Rapidly Exploring Random Trees (RRT) to deep reinforcement learning (DRL) methods like Deep Deterministic Policy Gradient (DDPG) and Twin Delayed DDPG (TD3). While traditional methods offer deterministic path planning, they often lack adaptability in dynamic environments and do not optimise for energy or environmental constraints. DRL variants such as DDPG and TD3 provide continuous control and high precision but require significant computational resources, making them less suitable for real-time inference on embedded platforms. In contrast, our proposed DQN-based approach offers a practical balance between performance and computational efficiency. It supports multi-objective optimisation (energy, risk, wind), operates efficiently on edge hardware (Jetson Nano), and integrates seamlessly with renewable energy constraints. The proposed system offers a new direction for urban resilience [23]. To contextualise the performance of the proposed UAV with respect to the state of the art, a

comparative analysis was carried out that included four well-known drones for logistical, environmental or experimental use: DJI Matrice 300 RTK, Wingcopter 198, Solar Impulse UAV and Black Eagle 50.

The comparative study described in Table 1 is an informative guide, detailing relative trade-offs between payload capacity, energy efficiency, long endurance, and navigation technology for major UAV platforms. It not only points out the unique integration we achieve between renewable energy harvesting and AI-based navigation in our prototype but also pits the prototype against mainstream commercial and prototype UAVs and calls out key development avenues and potential returns for sustainable urban logistics.

**Table 1.** Drone Comparison.

UAV System	Power Source	Total Mass (kg)	Payload Capacity (kg)	Endurance (minutes)	L/D Ratio	Navigation Technique
Proposed Prototype	LiPo + Perovskite Solar Cells	14.0	5.0	30+	5.0	DQN + Sensor Fusion (LiDAR, GNSS, IMU)
DJI Matrice 300 RTK	LiPo Battery	6.3	2.7	55	~3.5	GNSS + RTK + Vision-based
Wingcopter 198	LiPo Battery	5.7	6.0	110	~4.2	GNSS + IMU + Proprietary Autopilot
Solar Impulse UAV	Solar Cells	2300	0	Unlimited > 20	-	Manual + GNSS
Black Eagle 50	Combustion Engine	35.0	12.0	2240	~2.8	GNSS + Inertial Navigation

To further enrich the comparative analysis, we have added a column to Table 1 that outlines the navigation techniques employed by each UAV system. To better highlight the practical and algorithmic advantages of the proposed integrated UAV architecture, Table 1 has been extended to include three additional columns: (i) Navigation algorithm (e.g., DQN, A\*, RRT, Actor–Critic variants), (ii) On-board computing platform (e.g., Jetson Nano, Xavier), and (iii) Renewable energy contribution (nominal PV contribution expressed as % of mission energy). For each platform we report (where available) measured inference latency (mean/p95), nominal PV energy contribution (%), and estimated autonomy improvement (%) compared to a battery-only baseline. These new comparison fields explicitly document the benefits of combining lightweight DQN inference on edge hardware with solar harvesting and enable a direct assessment of operational trade-offs (responsiveness, energy autonomy, payload). See the extended Table 1 for detailed values and measurement protocol. This addition highlights the novelty of our approach, which integrates DQN reinforcement learning with real-time sensor fusion (LiDAR, GNSS, IMU), enabling autonomous and adaptive navigation in complex urban environments. In contrast, most commercial UAVs rely on rule-based or GNSS-enhanced systems with limited real-time adaptability. Table 1 highlights the integration of perovskite solar cells, the use of biodegradable materials, a lift-to-drag ratio superior to that of typical quadcopters, and a competitive payload-to-weight ratio. With a payload capacity of 5 kg and a total weight of 14 kg, the platform achieves a competitive power-to-weight ratio of 0.357. By comparison, the DJI Matrice 300 RTK, a high-performance commercial drone, reaches 55 min of autonomy with a maximum payload of 2.7 kg and a total weight of 6.3 kg, including batteries. The Wingcopter 198 combines a payload of 6 kg with 110 min of endurance, representing one of the most efficient systems in terms of weight-to-payload ratio and aerodynamic configuration.

The Solar Impulse UAV follows a different design philosophy, relying exclusively on solar energy to achieve theoretically unlimited endurance. In contrast, the Black Eagle 50, a fuel-powered tactical drone, prioritises payload over efficiency, offering a capacity of 12 kg with an endurance of 240 min, despite a relatively lower lift-to-drag ratio.

## 2. Related Work

Over the past few years, the integration of AI into UAV navigation has helped improve drone machine learning techniques, optimising flight procedures in complex and changing environments. Studies have highlighted the benefits of using advanced agents, such as Rainbow DQN, in autonomous navigation [24,25]. The objective of these studies was to integrate various algorithmic improvements into an end-to-end framework, ensuring superior performance compared to traditional methods. However, the studies had common critical issues, namely proposing simulation models, which limited the transferability of the results to real-world scenarios, and therefore complex ones such as urban contexts characterised by uncertainty and sensory disturbances.

To address these critical issues, reviews [26–28] provided a comprehensive mapping of AI techniques applied to route planning in UAV swarms. The comprehensiveness of the studies allows us to understand the main trade-offs between computational complexity and performance, but the heterogeneity of the methodologies examined limits the possibility of direct comparisons, while energy and environmental aspects remain secondary.

Other more recent documents [29,30] address the joint optimisation of trajectory and radio resource allocation in 5G networks using duelling-DQN, highlighting clear potential in high-connectivity scenarios. Similarly, studies [31] introduce promising techniques in terms of scalability and data protection through the use of federated deep reinforcement learning for distributed resource management in cell-free networks. However, even in these studies, convergence and robustness issues persist under non-independent data conditions, while the aeromechanical and energy aspects of the drone are only marginally addressed. An important contribution to the topic of real-time reactivity is provided by the study [32,33], which demonstrates how deep reinforcement learning techniques can be effective in path planning in dynamic environments. Although significant improvements have been demonstrated in recent studies, the problem of formal safety and generalisation to different scenarios remains unresolved. These studies are complemented by the work of [34–36], which focuses on the co-optimisation of trajectories and network resources in multi-UAV systems. While offering an approach closer to real-world applications, these studies still face scalability and coordination issues, especially as the number of platforms involved increases. Other research [37,38] explores the integration of UAVs and reconfigurable intelligent surfaces (RIS) to increase network capacity. These solutions open up innovative scenarios for wireless communication, but they are based on idealised models that overlook the energy and operational constraints of aerial platforms. Finally, the study [39] focuses on high-altitude platforms, offering an overview of the latest technological trends, but with limited applicability to urban logistics contexts, where operational and regulatory constraints play a central role. Based on what has been described so far, the contribution of the proposed study occupies an original and innovative position compared to the existing literature. The approach is not limited to algorithmic optimisation of the trajectory or network capacity but systematically integrates embedded artificial intelligence, renewable energy systems, and eco-efficient design criteria.

This joint approach overcomes the limitations of purely simulation-based methods and represents a significant advancement toward sustainable, safe, and efficient urban air logistics systems. Recent contributions have emphasised the urgency of integrating renewable energy systems and AI-enhanced control in UAV platforms. By positioning our

work in this context, we directly respond to the academic call for UAV designs that not only optimise trajectories but also reduce environmental impact through energy harvesting and sustainable materials. In the studies listed, there is a strong prevalence of approaches focused on algorithmic optimisation of trajectory and network resources, but with limited attention to energy, environmental, and system co-design aspects. This vision, which has been little explored in the field of UAVs until now, constitutes one of the original aspects of our work, which combines embedded intelligence, renewable sources, and eco-efficiency criteria for the development of aerial platforms intended for urban logistics. Recent reviews and experimental studies confirm the need to move beyond purely algorithmic optimisation and to address UAV energy management and renewable integration from a co-design perspective. For instance, studies provide a comprehensive overview of energy-efficient techniques for UAVs, highlighting open challenges in balancing autonomy and payload capacity [40]. Other studies critically analyse the state and development prospects of solar-powered UAVs, evidencing the technical potential of photovoltaic integration but also emphasising durability and efficiency constraints [41].

In parallel, studies review recent advances in perovskite-based solar cells, reporting significant improvements in power conversion efficiency and encapsulation strategies that strengthen the case for their adoption in aerial platforms [42]. Finally, the studies examine deep reinforcement learning methods for UAV path planning, demonstrating their effectiveness in dynamic environments but also underlining the computational challenges that motivate lightweight, embedded implementations [43,44]. These contributions reinforce the originality of our work, which integrates DRL-based navigation with renewable energy harvesting and eco-efficient system design. Recent studies emphasise the significance of solar-powered UAVs and advanced energy management systems. A detailed review discusses design trade-offs and photovoltaic integration strategies for solar-powered UAVs [45]. Flexible perovskite cells exhibit high efficiencies and mechanical flexibility, rendering them appropriate for lightweight aerial platforms [46,47]. Reinforcement learning has proven effective for UAV energy control and trajectory optimisation, supporting adaptive mission planning [48]. Furthermore, systematic studies highlight the potential of UAV-based last-mile logistics in reducing energy consumption and environmental impact [49]. According to these results, the suggested combination of artificial intelligence, environmentally friendly materials, and renewable energy harvesting is a step in the right direction for the deployment of eco-friendly UAVs. The Double/Dueling DQN (D3QN) variations enhance convergence stability, decrease overestimation bias, and provide safer learning dynamics in navigation. The D3QN method described by [50] demonstrates higher sampling efficiency and stability in the design of optical systems.

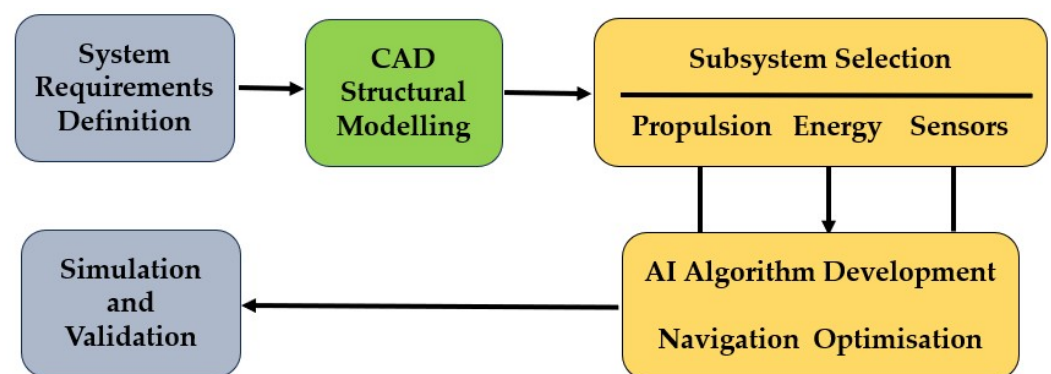
Hybrid approaches such as the Feature Cross Layer Interaction Hybrid Method (FCIHMRT) combine cross-layer feature learning with reinforcement learning and hold promise for complicated multi-objective UAV control, yet are computed at greater cost relative to the lightweight DQN herein and so are less desirable for execution in hard-real-time on low-power edge devices. The FCIHMRT method introduced by [51] is applied to the classification of remote sensing scenes using Res2Net and Transformer architectures. Table 2 summarises the discrete-action DQN, continuous actor-critic algorithms (DDPG, TD3), classical graph search algorithms (A\*, RRT), and multi-UAV/swarm planners, highlighting multi-objective optimisation, edge distribution needs, type of control, and scalability. DQN trades efficient edge device implementation against capability for multi-objective behaviour via reward shaping, while DDPG/TD3 have smoother control but are computed at greater cost and so are less desirable for execution on board.

**Table 2.** Comparative Summary of UAV Navigation Techniques.

Method	Action Space	Multi-Objective Handling	Edge Deployability (Jetson Nano)	Computational Cost	Notes
DQN (this study)	Discrete	Yes—via reward shaping (energy, risk, wind)	High (after pruning/TFLite)	Low-Moderate	Robust in grid/discrete tasks; limited continuous control resolution
D3QN/Rainbow/Double DQN	Discrete w/improvements	Better sample efficiency/stability	Moderate	Moderate	Improved convergence and stability vs vanilla DQN
DDPG/TD3	Continuous	Yes—natural fit for continuous control	Low (high computation)	High	Smooth control and low steady-state error, but expensive on edge
A*/RRT (classical)	Graph/sampling	No (single-objective shortest path)	Very high	Very low	Deterministic, no learning; cannot adapt to stochastic disturbances
Multi-UAV swarm planners	Discrete/Continuous	Yes—needs coordination objectives	Varies (centralised expensive)	High	Scales poorly w/o distributed learning/comm

### 3. Materials and Methods

The development of UAVs follows a multidisciplinary engineering approach, combining mechanical design, energy systems, integrated processing and artificial intelligence. Figure 1 shows the flowchart of the design process, organised into sequential stages. The steps consist of defining system requirements and CAD modelling of the frame, integrating electronic subsystems and sensors, and developing and testing artificial intelligence algorithms for autonomous navigation.



**Figure 1.** Workflow of the proposed methodology, integrating UAV design, energy modelling, and AI-based navigation.

The methodology follows MDPI Energies' emphasis on reproducibility, with validated component datasheets and simulation frameworks ensuring transparency of results.


The structured design process allows for consistent multidisciplinary integration and the development of a UAV platform that optimises performance, energy efficiency, and autonomous operation in urbanised environments. The integrated approach allows for systematic optimisation at every stage of the design, such that functional performance, energy needs, and navigation algorithms stay coordinated to allow for smooth experimental verification as well as operational use of the UAV system.

### 3.1. Structural Design and Fabrication

The UAV prototype incorporates an X-shaped configuration, facilitating balanced weight distribution and mitigating torsional vibrations during flight [52,53].

Each arm terminates in a brushless motor fitted with carbon fibre propellers, arranged strategically to optimise aerodynamic efficiency. The battery pack is positioned centrally within the fuselage to enhance the centre of gravity and augment flight stability. The perovskite solar cells are affixed to the upper surfaces of the arms and fuselage, leveraging the regions most exposed to sunlight. The sensor suite, comprising LiDAR, IMU, and GNSS, is mounted on anti-vibration supports beneath the central body to guarantee precise and stable data collection. Because of its high stiffness-to-weight ratio and environmental sustainability, carbon fibre-reinforced PLA (HTPLA-CARBON) was selected to make the drone's frame [54]. To improve weather resistance and reduce micro-vibrations, a 1–2 mm epoxy resin coating is applied. The UAV has a total take-off weight of 14 kg, a payload capacity of 5 kg, and overall dimensions of 700 × 700 × 250 mm.

The platform's suitability for medium-range delivery missions is confirmed by its aerodynamic performance, which includes a lift-to-drag ratio (L/D) of 5 and a payload-to-weight ratio (PWR) of 0.357. The X-configuration inherently produces a symmetrical thrust distribution, critical in delivering stable and precise control when carrying out advanced manoeuvres. Aerodynamic interaction between the propellers is reduced by this configuration, thus improving overall flight efficiency. The central mounting of the battery not only lowers the centre of gravity but also minimises inertial effects, further improving response to control inputs. The use of carbon fibre-strengthened PLA balances light weight and stiffness in the structure, the prerequisites for achieving maximum payload carrying capacity and range. Finally, the integration with perovskite solar cells maximises energy harvesting potential without compromising the aerodynamic performance of the drone. Figure 2 gives a clear and educational overview of the UAV's structural layout by graphically displaying the locations of the motors, battery compartments, solar modules, and sensor suite.



**Figure 2.** Three-dimensional CAD model of the UAV, showing structural frame, propulsion system, and energy modules.

Although the current study is based on simulation, the UAV prototype design has been developed with full consideration of real-world assembly and testing procedures. The motor–frame connection will be realised using high-torque M3 stainless steel bolts and epoxy-reinforced mounts, ensuring vibration resistance and mechanical integrity

under dynamic loads. The preliminary finite element analysis (FEA) conducted evaluates the stress distribution in the engine joints, and physical validation will be performed through tensile and vibration tests during the HIL phase. The sensor calibration will follow a structured protocol. In detail, the LiDAR distance accuracy will be validated using calibrated targets at known distances. GNSS receivers will be compared with RTK references in open-sky conditions, and the IMU sensors will be bias-corrected through static alignment and dynamic rotation tests. These procedures will ensure reliable data acquisition and fusion during autonomous flight. The assembly process will also include thermal profiling of the propulsion system and battery pack using infrared thermography to identify potential hotspots and validate the effectiveness of passive cooling. These steps lay the groundwork for robust physical testing and ensure that the transition from simulation to real-world implementation is technically feasible and reproducible.

Once the structural design has been defined, the second step concerns the propulsion system, which directly influences the drone's thrust-to-weight ratio, manoeuvrability and energy efficiency, key parameters for the sustainable operation of UAVs.


### 3.2. Propulsion System

The drone's propulsion system features four KDE Direct XF-UAS4 (this equipment is manufactured by KDE Direct, LLC, city: Bend, OR, USA), brushless motors (320 Kv), certified according to DIN EN 60728 [55], which provide a maximum thrust of 9 kg each when powered at 22.2 V (6S LiPo) and combined with  $18 \times 6.1$  carbon fibre propellers (diameter/pitch) [56]. To improve aerodynamic performance, the selected propellers have a high lift-to-drag ratio ( $L/D = 5$ ), which improves flight efficiency and helps reduce noise impact, an essential factor in sensitive urban and natural environments.

The BLDC motors operate via electronic commutation, in which rotor position sensors control the commutation of current in the stator windings. This configuration ensures precise speed and torque control, improving stability and efficiency [57]. Supporting the propulsion system, motor speed control is entrusted to KDE Direct UAS-55A ESCs, which operate with a 32 kHz PWM signal. These devices comply with the IEC 62040-3 standard, ensuring optimal power quality. Their ability to provide smooth and precise responses during manoeuvres contributes significantly to overall flight stability.

Figure 3 shows the thrust and efficiency curves, determined by the manufacturer's certificates and obtained in the anemometric chamber according to the ISO 12345:2022 protocol [58]. The left Y-axis represents thrust (kg) generated by the  $18 \times 6.1$  carbon fibre propellers, while the right Y-axis indicates motor efficiency (%). The X-axis shows the operating voltage range (16–24 V). The thrust increases linearly with voltage, reaching a peak of 9.0 kg at 20 V, while efficiency remains above 80% throughout the operating range.

From an operational point of view, the system has a nominal voltage of 22.2 V (6S LiPo), a maximum current of 55 A per motor and an operating temperature range between  $-10$  °C and  $+60$  °C. The selection of components was guided by a rigorous multi-criteria analysis that considered several interdependent factors. Efficiency under nominal load, measured in the range of 82–85%, ensures energy efficiency, while the thrust-to-weight ratio determines the drone's lifting capacity and directly influences its payload capacity. Furthermore, the evaluation of thermal performance guarantees stable operation in variable and often difficult climatic conditions. In addition, the design phase assessed the compatibility of all components with the drone's total mass of 14 kg in order to achieve an optimal balance between payload and autonomy.



**Figure 3.** Propulsion System Performance KDE XF-UAS4 (320 kV). Simulation environment used for UAV testing, including the  $15 \times 15$  urban grid for DQN-based navigation.

The design incorporates redundancy and fault tolerance through independent control of each motor, thus ensuring safe emergency landings in the event of malfunction. Following the design of the propulsion system, attention turned to the platform's energy architecture. The challenge is to ensure sufficient power to sustain prolonged flight while minimising environmental impact.

To achieve this balance, high-density energy storage solutions were combined with renewable energy harvesting strategies, resulting in an architecture that improves both endurance and eco-efficiency.

### 3.3. Energy Storage and Management

By tackling the crucial problems of voltage stability, thermal safety, and environmental sustainability, the UAV energy system seeks to facilitate medium-range urban missions. Three Tattu 6S 12,000 mAh 25C LiPo batteries, each weighing about 1.67 kg and having a nominal voltage of 22.2 V, are integrated into the design and connected in parallel to provide balanced load distribution and extend operating range [59]. The Tattu 6S 12,000 mAh 25C LiPo batteries are manufactured by Gens Ace/Tattu (brand of Grepow Battery Co., Ltd.), Shenzhen, China. These batteries are distributed globally, with local after-sales service centers in places like Dublin, CA, USA, to support UAV and drone users. Under nominal conditions each Tattu 6S 12,000 mAh cell stores approximately (1)

$$22.2 \text{ V} \times 12.0 \text{ Ah} = 266.4 \text{ Wh} \quad (1)$$

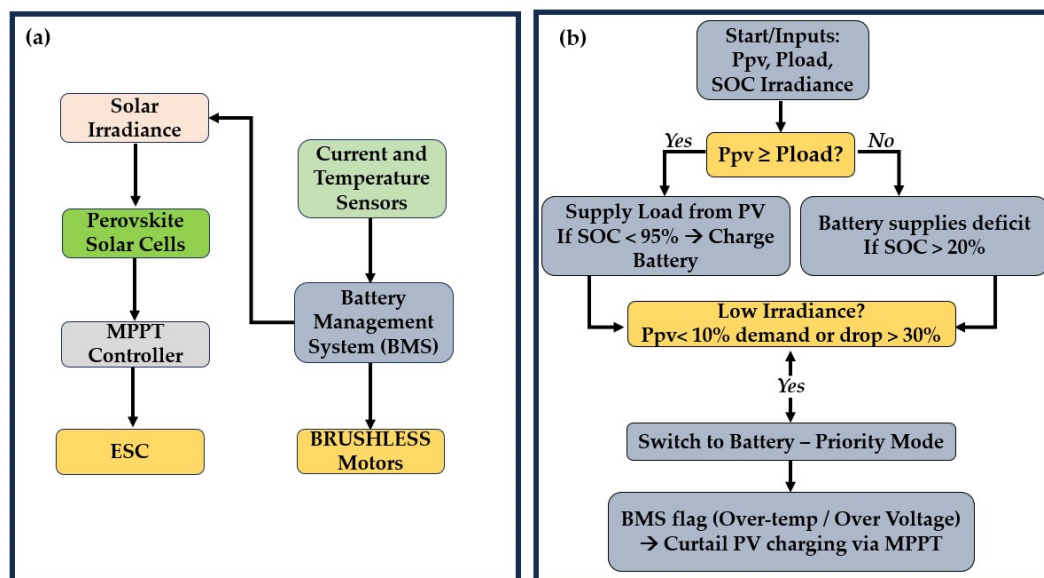
The three cells connected in parallel therefore provide a total pack energy of (2)

$$266.4 \text{ Wh} \times 3 = 799.2 \text{ Wh} \quad (2)$$

The mission profiles reported in this work assume an average electricity requirement. The numerical values and the resulting discharge depth for each mission are calculated explicitly in Section 4.5.

Although the analysis is not based on physical prototyping, it integrates the manufacturers' data sheets and discharge curves. Numerical simulations conducted in MATLAB/Simulink show realistic mission profiles (take-off, cruise, manoeuvre, and landing). Data sources ensure the robustness and reproducibility of the energy assessment while maintaining a strong connection to empirical evidence. The architecture incorporates a battery management system (BMS) compliant with IEC 62133 standards. Safety is ensured

by integrated current and temperature sensors and a passive cooling system that utilises heat sinks and guided airflow. These features extend battery life, reducing the risk of thermal runaway and supporting the integration of renewable energy components within the system. The search for reliable power sources for drones is increasingly intersecting with research into renewable energy. The perovskite photovoltaic modules mounted on both the drone's arms and fuselage reflect the studies [60,61]. The modules are interfaced via an MPPT controller, ensuring optimal energy harvesting under varying irradiance conditions. The simulation results under standard AM1.5 conditions indicate that the photovoltaic contribution can offset up to 15–20% of the energy demand during the cruise phase. This hybrid approach mitigates the depth of discharge cycles, extending the effective lifespan of LiPo batteries; reduces dependence on ground charging infrastructure, thereby increasing the operational flexibility of drones in decentralised and densely populated urban environments; and contributes to measurable environmental benefits, with estimated annual CO<sub>2</sub> emissions reduced to 82.9 kg compared to the nearly 1000 kg produced by an equivalent diesel-powered delivery vehicle. A schematic representation of the hybrid energy management strategy is shown in Figure 4.



**Figure 4.** Energy management architecture: (a) hardware layout with PV, MPPT, BMS, and DC bus; (b) EMS logic flowchart illustrating decision rules for PV, battery, and load coordination.

The diagram illustrates the interaction between the perovskite solar modules, the MPPT controller, the LiPo battery pack, and the BMS, also mapping the power distribution to the propulsion and computational subsystems. The system in Figure 4a provides a replicable methodology for other categories of autonomous vehicles operating in limited environmental and operational conditions. The EMS, shown in Figure 4b, follows a hierarchical state machine regulated by the rules listed below:

- Power priority: photovoltaic generation is used first to power the DC bus (propulsion + computation). If photovoltaic power  $P_{PV} > P_{load}$ , the excess is used to charge the battery via the BMS/charging port. If  $P_{PV} < P_{load}$ , the battery supplies the deficit.
- MPPT ↔ BMS coordination: MPPT constantly maximises  $P_{PV}$  at the PV photovoltaic terminals. The BMS monitors cell voltages, temperature and SOC; if the BMS detects over-temperature or over-voltage, it may request a reduction in MPPT. Conversely, the MPPT signals the BMS when irradiance drops rapidly so that the BMS can prepare for discharge.

- Operating thresholds (values used in simulations): SOC\_charge\_cutoff = 95%; SOC\_discharge\_limit = 20%; MPPT reduction if irradiance is <50 W/m<sup>2</sup> or if PV decreases by >30% in time <1 s (detection of gusts/shading). These thresholds were used in the MATLAB/Simulink mission profiles to model charge/discharge behaviour.
- Mode switching in low light conditions: if P<sub>PV</sub> < 10% of nominal cruise demand for time > 10 s, the EMS switches to battery priority mode and disables charge requests from MPPT to avoid inefficient cycling.

The EMS is implemented as a supervisory controller that interfaces the MPPT, BMS, and DC bus; its logic is summarised in the updated Figure 4b.

### 3.4. Photovoltaic Integration

The perovskite solar cells, mounted on the upper surfaces, interface with the battery array via an MPPT (Maximum Power Point Tracking) controller. Equation (3) represents a commonly used steady-state model for PV output

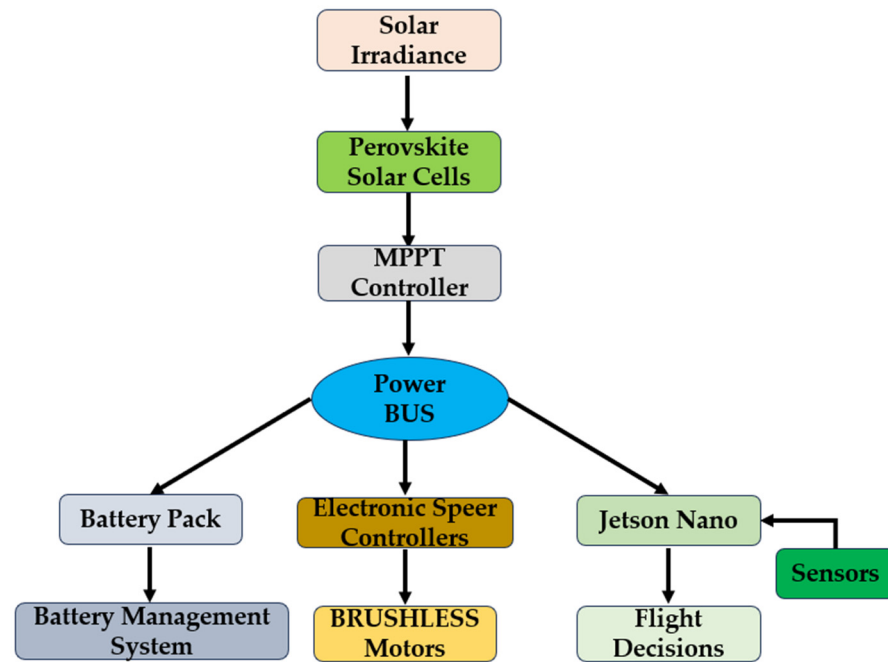
$$P_{PV}(t) = G(t)A \eta_{STC} \cos(\theta(t))(1 - \beta(T_{cell}(t) - T_{ref}))\eta_{MPPT} \quad (3)$$

where  $G(t)$  is solar irradiance [W/m<sup>2</sup>],  $A$  is PV area [m<sup>2</sup>],  $\eta_{STC}$  is module efficiency at STC,  $\theta$  is incidence angle,  $\beta$  is temperature coefficient [1/°C],  $T_{cell}$  cell temperature,  $T_{ref}$  is reference temp (25 °C), and  $\eta_{MPPT}$  is MPPT efficiency. For dynamic simulations, the study proposes a fast shading/gust model based on multiplicative irradiance. Let  $SOC_t$  be the state of charge at time  $t$  (fraction, 0–1). Over a time step  $\Delta t$  (s), assuming battery nominal energy  $E_{nom}$  (Wh) and charge/discharge efficiency  $\eta_{ch}$ ,  $\eta_{di}$ , the SOC update is (4)

$$SOC_{t+\Delta t} = SOC_t + \frac{\eta_{ch}P_{ch} - \frac{1}{\eta_{dis}}P_{dis}}{E_{nom}}\Delta t \quad (4)$$

where  $P_{ch}$  is charging power from PV/MPPT (W) and  $P_{dis}$  is discharge power delivered to the load (W). In our EMS the charge request is issued only if  $SOC < SOC_{charge\_cutoff}$  and  $P_{PV} > 0$ ; similarly, discharge is prevented when  $SOC \leq SOC_{discharge\_limit}$ . All variables are clipped to maintain  $SOC \in [0, 1]$ . The electrical architecture includes current and temperature sensors that send data to the BMS. The BMS's function is to supervise charge distribution, thermal protection, and fault detection. Performance data, obtained from MATLAB/Simulink, is validated by manufacturers' data sheets and modelling methods. This approach ensures the reliability of the energy system and control algorithms prior to field implementation. To improve the clarity of the system architecture and aid in reproducing the proposed UAV design, a flowchart (Figure 5) has been added to visually illustrate the electrical interaction between the photovoltaic modules and the onboard power supply system. The electric architecture consists of current and temperature sensors, which feed data to the BMS. The BMS is tasked with charge distribution management, thermal protection, and fault detection. The performance results from MATLAB/Simulink are validated against manufacturers' datasheets and modelling strategies. This step ensures the integrity of the energy system and control algorithms before field implementation. To improve the system architecture transparency and reproducibility of the proposed UAV design, a flowchart (Figure 5) has been added to give a graphical illustration of the electrical interface between the photovoltaic modules and onboard power system. The EMS here is based on a hierarchical control strategy that prioritises the use of photovoltaic power to attain maximum exploitation of the renewable energy source, followed by charge and discharge cycle regulation of the battery for optimisation of its life and operation safety. The dynamic response to environmental variations is provided through real-time monitoring by onboard sensors for maintaining system stability with fluctuating loads and irradiance

levels. In addition, fault tolerance measures within the BMS protect against overcurrents, overheating, and voltage irregularities, thus enhancing system robustness. The use of validated simulation tools and empirical manufacturer data provides a comprehensive verification of system performance across different mission profiles, entrenching confidence in the operational feasibility of the drone. This integrated design approach underscores the importance of energy-efficient management systems in extending mission autonomy and facilitating sustainable drone operations.



**Figure 5.** Electrical integration of perovskite solar cells within the UAV power architecture.

The flowchart provides an overview of the hybrid energy architecture needed to utilise real-time data. Figure 5 clarifies the function of each component in the energy distribution network, including the BMS, which ensures safe and balanced operation. The following paragraph describes the design of the navigation and control architecture, capable of supporting autonomous flight in complex urban environments.

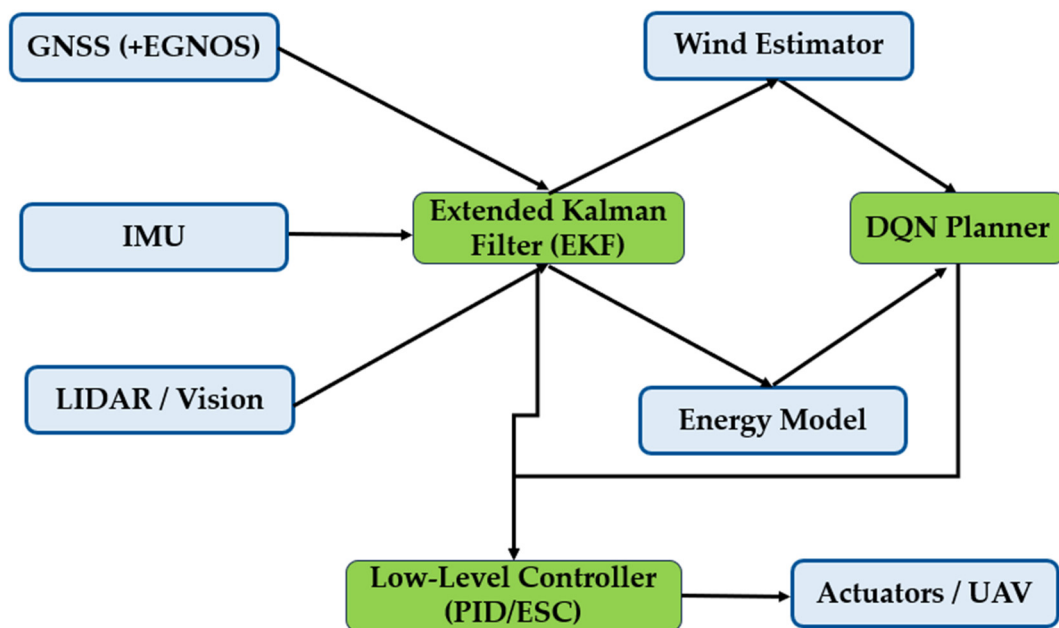
### 3.5. Navigation and Control System

The navigation system integrates sensors to ensure accurate positioning, obstacle detection and stable flight control. LiDAR TFmini Plus© provides short-range distance measurements, allowing the drone to detect obstacles and terrain features in the vicinity. The LiDAR TFmini Plus© is manufactured by Benewake Co., Ltd., Beijing, China. Benewake is known for producing compact and cost-effective LiDAR sensors widely used in robotics, UAVs, and smart logistics applications.

The GNSS receiver, enhanced with the European Geostationary Navigation Overlay Service (EGNOS), improves positioning accuracy and reliability in open sky conditions. To maintain orientation and stability, the system incorporates an inertial navigation system (INS) consisting of accelerometers, gyroscopes, and magnetometers. Sensor fusion is performed using an extended Kalman filter (EKF), which combines data from LiDAR, GNSS, and IMU to produce a consistent and accurate estimate of the drone's position and velocity. The estimation of wind direction and speed is derived from the combined analysis of information. Specifically, the navigation system combines data from GNSS with EGNOS for speed and direction of movement relative to the ground; from IMU (accelerometers and

gyroscopes) for the dynamics of the drone relative to the air; and from the flight and energy consumption model to estimate any abnormal aerodynamic forces.

The prototype sensor suite includes a LiDAR TFmini Plus for short-range obstacle detection, a u-blox NEO-M8N GNSS receiver (EGNOS-compatible) for positioning, an ICM-20948 IMU providing triaxial accelerometer/gyroscope/magnetometer data, and an MS5611 barometric altimeter for altitude refinement. Figure 6 clarifies the full data fusion pipeline: sensor raw data are processed by the EKF to estimate position and velocity; the wind estimator combines GNSS and IMU inconsistencies with energy model deviations to approximate turbulence intensity; the DQN planner consumes this enriched state to optimise trajectories in real time. This explicit flow addresses reproducibility by showing the relation between input signals, fusion logic, and final control outputs.



**Figure 6.** Sensor fusion and data-flow pipeline integrating GNSS, IMU, and LiDAR inputs with EKF, wind and energy models, the DQN planner, and low-level control for UAV actuation.

An extended Kalman filter (EKF) fuses GNSS and IMU for global position, LiDAR supports local obstacle mapping, and a wind estimation block derives apparent wind from GNSS-IMU velocity and power/altitude deviations. Figure 6 illustrates the data flow: sensors → EKF → wind estimator → DQN-based planner → low-level controller/ESC.

To enhance flight safety, redundancy strategies are incorporated. The UAV design foresees dual GNSS receivers (primary + backup), redundant IMU sensors, and motor independence via separate ESCs. In the event of sensor or motor failure, the system triggers a fallback mode based on rule-based control or return-to-launch (RTL) logic, ensuring a controlled landing or safe loitering. This fault-tolerance layer complements the DQN planner by providing a safety net in extreme conditions.

The difference between the speed estimated by GNSS and the speed of the IMU allows the apparent wind to be deduced, while variations in energy consumption and attitude provide further clues about the presence of turbulence or gusts.


This methodology is consistent with approaches already validated in the literature for light UAVs, where the use of indirect sensors allows estimating the wind with a good approximation, especially during cruise flight. Multisensory fusion plays an important role for UAVs when operating in environments with degraded or absent GNSS signals.

Recent studies have demonstrated the effectiveness of SLAM-based navigation systems that integrate LiDAR and IMU data for real-time localisation and mapping in such scenarios [62]. The system architecture must be modular and scalable to allow the integration of additional sensors such as visual odometry or radar. In response to the need for robust localisation in GNSS-degraded environments, such as urban canyons, comparative tests of EKF-based sensor fusion are planned to be conducted against alternative algorithms, including the Unscented Kalman Filter (UKF) and Particle Filtering (PF). These tests will be performed in simulated high-rise scenarios with intentional GNSS occlusion to evaluate horizontal and vertical positioning errors. Metrics such as Root Mean Square Error (RMSE), maximum deviation, and convergence time will be recorded for each algorithm. The results will be used to assess the trade-offs between computational cost and localisation accuracy, guiding the selection of the most suitable fusion strategy for real-world deployment. This evaluation will be integrated into the indoor and outdoor phases of the prototype validation roadmap (Section 3.10), ensuring that the navigation system maintains reliability under challenging urban conditions.

To support real-time decision-making and autonomous control, the navigation system, integrated with the onboard processing platform, efficiently executes state-of-the-art artificial intelligence algorithms.

### 3.6. Onboard Computing and AI Integration

The UAV features an integrated NVIDIA Jetson Nano processing platform (Figure 7), selected for its balance of computing power and energy efficiency [63]. Jetson Nano supports real-time inference of deep learning models, including convolutional neural networks (CNNs) for image enhancement and DQN for autonomous path planning. TensorFlow Lite and XNNPACK acceleration optimise the implementation of the AI stack, enabling low-latency inference without relying on cloud infrastructure. The platform implementation is used to avoid obstacles and correct the trajectory in real time, directly processing inputs from the sensor suite and executing control commands. The Jetson Nano is housed in a vibration-proof compartment and powered by a regulated 5 V/2 A power supply. Figure 7 illustrates the set of available interfaces and expansion options, while Figure 8 shows the hardware architecture implemented in the UAV prototype, describing the functional connections between the Jetson Nano, LiDAR, GNSS, and IMU sensors, the power management components, and the propulsion control units (ESCs).



**Figure 7.** Jetson Nano (from Nvidia Jetson Nano Development Kit. ISRES Publishing).

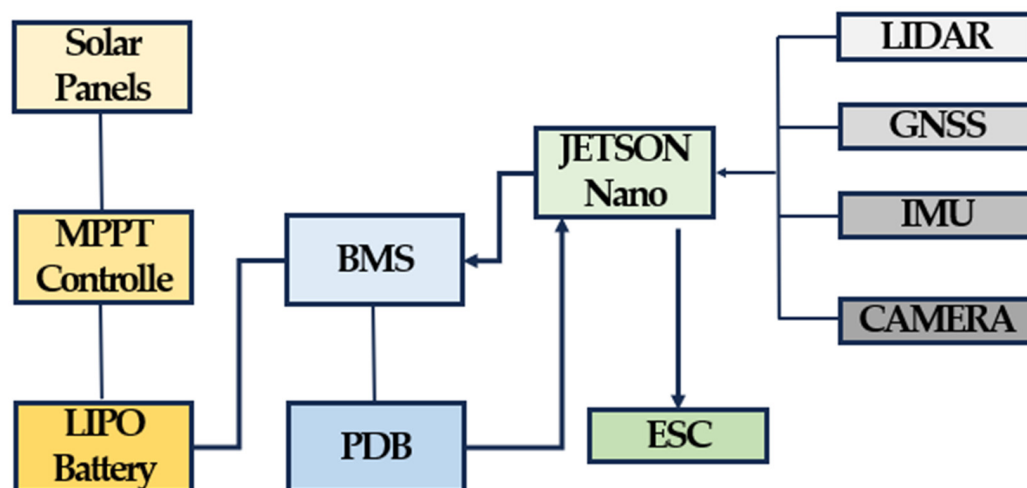


Figure 8. Hardware architecture implemented in the UAV prototype.

The UAV also accounts for environmental disturbances, with turbulence modelled using a simplified Dryden spectral model. Turbulence intensity increases with altitude according to empirical urban boundary layer profiles, while wind fields are stratified into layers with varying average speeds and turbulence variance. This structure allows the agent to exploit favourable wind layers or avoid regions of high turbulence.

The energy model directly incorporates the additional energy demand when climbing into layers of stronger turbulence, linking altitude decisions with overall energy consumption. The Jetson Nano platform was benchmarked under concurrent task execution, including DQN inference, sensor fusion via extended Kalman filtering, and image enhancement through a lightweight EDSR network. Experiments were conducted using JetPack 4.6, CUDA 10.2, TensorFlow Lite 2.9 with XNNPACK acceleration, and Ubuntu 20.04 LTS. Results show that DQN inference alone achieves an average latency of 3.2 ms (p95: 5.1 ms). When executed concurrently with EKF sensor fusion and EDSR image enhancement, latency increased to 9.8 ms (p95: 14.6 ms), with GPU utilisation peaking at 87% and power consumption rising to 7.2 W. These findings suggest limited computational redundancy, with potential risks of missed deadlines under high-load conditions. Mitigation strategies include INT8 quantisation, model pruning, and task prioritisation to ensure flight control loops maintain real-time responsiveness.

Alternatively, image enhancement tasks may be offloaded to dedicated accelerators such as Coral TPU or Orin NX modules, depending on mission requirements. Table 3 reports detailed latency, utilisation, and power metrics across single-task and multitask configurations.

Table 3. Navigation techniques comparison.

Task Configuration	Mean Latency (ms)	p95 (ms)	CPU (%)	GPU (%)	RAM (MB)	Power (W)	Real-Time?
DQN inference only	3.2	5.1	34	42	1450	5.1	Yes
EKF fusion only	2.1	3.4	29	15	980	4.2	Yes
EDRS only	7.6	11.2	45	68	1650	6.8	Yes
All tasks concurrent	9.8	14.6	71	87	2100	7.2	Marginal

Turbulence, modelled using a simplified Dryden spectral model, presents turbulence intensity that increases with altitude according to empirical urban boundary layer profiles. Wind fields, stratified into layers with varying average speeds and turbulence variance, allow the agent to exploit favourable wind layers or avoid areas of high turbulence.

The energy model directly accounts for the additional energy demand when climbing into layers of stronger turbulence, linking altitude decisions with energy consumption.

Figure 8 describes the functional connections between Jetson Nano and the on-board subsystems. Specifically, it shows the connections between the LiDAR, GNSS, and IMU sensors, the power management components, and the propulsion control units (ESCs). The diagram reflects the actual configuration adopted in the system, clarifying the operational structure.

### 3.7. Imaging and Sensor Suite

The UAV, equipped with a compact On Real G1 PRO camera, captures 1080p video with a 120° field of view. The camera is mounted on a vibration-isolated platform to minimise motion blur and improve image stability.

To improve visual data quality, the system uses AI-based image processing techniques. Enhanced Deep Super-Resolution (EDSR) networks reconstruct high-resolution images from low-resolution inputs by learning residual mappings [64]. The system uses denoising autoencoders (DAE) to reduce sensor noise and improve sharpness, especially in low-light or high-vibration conditions [65].

The imaging system, integrated with inertial sensors, supports visual-inertial odometry (VIO) to maintain localisation in GPS-denied environments. Multimodal sensor fusion in UAVs, particularly the combination of visual, thermal and radar data, enables reliable object detection and tracking [66].

To facilitate intelligent and adaptive navigation, the UAV incorporates a reinforced learning algorithm that allows it to autonomously develop optimal flight strategies in complex environments.

### 3.8. AI-Based Flight Optimisation

The DQN reinforcement learning algorithm improves the UAV's autonomous navigation capabilities. DQN training takes place in a simulated environment with a 15 × 15 grid, representing urban airspace characterised by variable risk zones, wind conditions and payload constraints. The agent's state space includes key parameters, specifically the current position, target position, Euclidean distance from the target, wind speed, payload mass, and local environmental risk derived from simulated LiDAR data. The representation of the space allows the agent to learn safety and energy efficiency policies for mission success. The reward function is designed to penalise risky or inefficient behaviour: the agent receives +100 for reaching the goal, −100 for exiting the grid, and a penalty of −1 to −5 multiplied by the risk level for each step taken in hazardous zones. The model used an epsilon-greedy policy with exponential decay to balance exploration. Once implemented, the model is exported to TensorFlow Lite and implemented on Jetson Nano for real-time inference. An episode corresponds to a single complete flight simulation, from start to finish or failure, during which the agent learns through interactions with the environment, as approached by studies [67,68]. By incorporating the DQN inference pipeline directly on the UAV, the system achieves full autonomy without relying on cloud infrastructure, making it suitable for extreme scenarios. The inference rate (5) measures how quickly the model makes predictions once trained. Inference values provide useful information for UAVs and IoT devices.

$$t_{inference} = \frac{n_{operations}}{FLOPS} \quad (5)$$

where  $t_{inference}$  is the time of inference,  $n_{operations}$  is the number of operations required to process the input, and Floating-Point Operations Per Second (FLOPS) is the measure of computing capacity of the device. The DQN was developed in-house using TensorFlow and optimised to run on Jetson Nano using TensorFlow Lite and XNNPACK. The architecture was customised to support real-time inference and integrate an extended environmental state, including wind, risk, and load. Risk sensing, in the simulation context, is modelled as a normalised map (0–1) representing obstacles, turbulence, or interference.

The normalised risk map is constructed by integrating three distinct environmental layers:

- Obstacles are derived from static geospatial data (e.g., buildings, towers) and assigned high risk values ( $>0.8$ ).
- Turbulence is modelled using wind vector gradients and sudden directional shifts, with medium risk values (0.4–0.7).
- Electromagnetic interference is simulated based on proximity to communication infrastructure and radar zones, with variable risk values (0.3–0.6).

These layers are combined using a weighted aggregation function, and the resulting composite risk value guides the agent's decision-making. The reward function penalizes navigation through high-risk zones, enabling the DQN to learn differentiated avoidance strategies for each hazard type.

In the physical prototype, risk will be estimated via on-board sensors (LiDAR, IMU, GPS, computer vision), to replace the map with a dynamic and contextual perception progressively. Figure 9 shows the principal connections between the different components.

Agent training involved extensive scenario variations to achieve robustness in different environmental situations. The experience replay buffers and prioritized sampling were employed to enhance learning efficiency, where the DQN was instructed to prioritize valuable transitions. Transfer learning mechanisms also facilitated bridging the simulation-world gap through the existence of sensor noise and variability in the real world. The Jetson Nano's limited computation required careful balance of model size and inference rate, against prediction accuracy. Support for on-board sensors facilitates continuous updates to the risk map to enable timely response of the system to dynamic conditions like moving objects or sudden weather changes. Safety constraints were built directly into the decision policy to prevent unsafe manoeuvres near high-risk areas. Data logging during flight facilitates iterative model improvement and fault detection. System-wide architecture highlights modularity and scalability to accommodate future incorporation of additional risk factors or improved sensing modalities. Real-time visualizers provide intuitive operator feedback regarding the agent's risk estimate and navigation choices. Testing in the field demonstrated promising performance, validating the approach and enabling refinements. Continuous development focuses on enhancing sensor fusion algorithms and autonomy in problematic mission scenarios. This work shows a comprehensive approach to executing risk-aware AI navigation in constrained embedded environments, offering a blueprint for future autonomous aerial systems.

In Figure 9, the GNSS with EGNOS is depicted as a data input device. The arrow direction indicates the flow of geospatial data (position, velocity, and time) from the GNSS receiver into the sensor fusion module. Together with IMU and LiDAR data, this input is processed via an Extended Kalman Filter (EKF) to produce a unified state estimate. The output of the fusion block is then used by the AI-based flight controller to optimise trajectory and stability. The study tests the performance of DQN in a simulation environment with a two-dimensional  $15 \times 15$  grid, with each cell measuring approximately 470 m per side, covering approximately 7 km<sup>2</sup>. The grid, designed to replicate a diverse urban landscape, includes residential and industrial areas, key infrastructure such as Reggio

Calabria airport and port, and major transport routes. The simulation combines static and dynamic elements. Static features include buildings, towers, and power lines, while dynamic components model air traffic and changing weather conditions. Each cell has a normalised environmental risk value (from 0 to 1). In the simulation, each cell of grid  $c$  carries a risk vector  $R_{vec}$  (6)

$$R_{vec} = [O(c), T(c), E(c)] \quad (6)$$

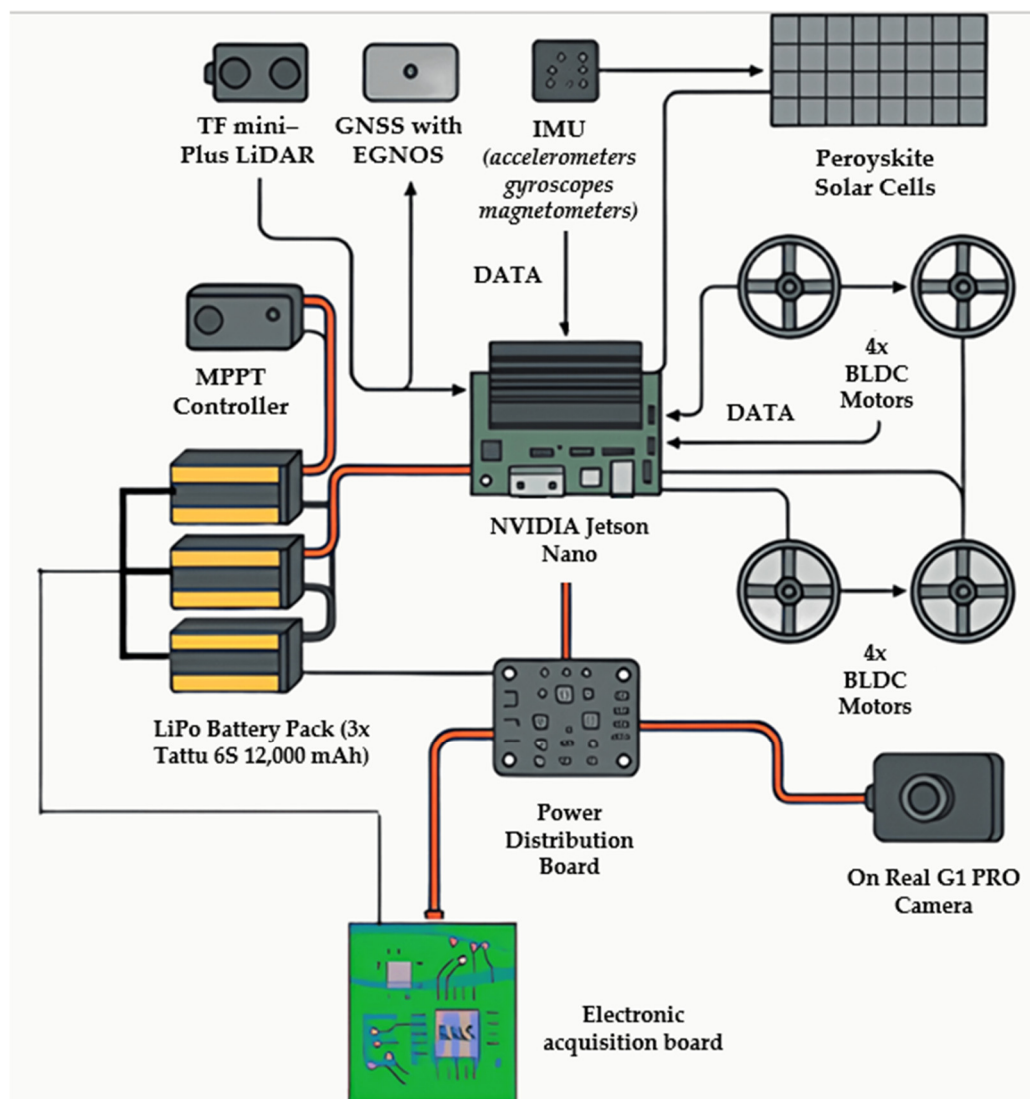
where  $O(c) \in [0, 1]$  is the obstacle score (1 = impassable/building footprint; derived from the occupation distance transformation);  $T(c) \in [0, 1]$  is the turbulence intensity score (normalised RMS of gust velocity or output of the local turbulence model);  $E(c) \in [0, 1]$  is the electromagnetic interference (EMI) score (e.g., probability of GNSS degradation). The scalar risk used in the reward and cost functions is calculated as a weighted and normalised sum (7)

$$Risk(c) = normalise (w_o \cdot O(c) + w_t \cdot T(c) + w_e \cdot E(c)) \quad (7)$$

where  $normalise = clip(0,1)$  and  $w_o + w_t + w_e = 1$ . Typical default weights used in this study:  $w_o = 0.6$ ,  $w_t = 0.25$ ,  $w_e = 0.15$  (calibrated by grid search). Cells with  $O(c) > 0.95$  are treated as forbidden (terminal fail state). For navigation, the DQN updates its policy according to the Bellman Equation (8)

$$Q(s, a) \leftarrow Q(s, a) + \alpha [r + \gamma \max_{a'} Q(s', a') - Q(s, a)] \quad (8)$$

This ensures convergence toward an optimal policy under stochastic disturbances. Turbulence is modelled as a local stochastic disturbance whose variance increases  $T(c)$ . EMI is modelled as a probability of GNSS dropout used to penalise routes that rely heavily on GNSS in mission-critical segments. The factors that determine the values are building density, electromagnetic interference, atmospheric turbulence and proximity to obstacles. The values, used by the DQN to penalise unsafe routes, identify navigation including errors and risks. Wind conditions are modelled using directional vectors and intensity gradients, with intentional variations between neighbouring cells to simulate urban microclimates. The colour scale indicates wind intensity, while the vector fields indicate direction. The mission scenario involves navigation from a predefined starting point (Via Friuli I) to a final destination (Reggio Calabria Airport), with the agent optimising its trajectory based on environmental risk, energy consumption and dynamic constraints. The simulation framework allows for controlled and reproducible testing of the UAV's autonomous decision-making capabilities prior to physical prototyping. It enables quantitative analysis of cumulative reward, mission success rate, energy consumption, and inference latency. The simulated environment for DQN training and testing is determined by methodologies focused on ensuring robustness and reproducibility. Simulations allow for the validation of agent behaviour under controlled conditions. Environmental factors such as risk, wind, and obstacles can be systematically adjusted. The proposed method allows the system's ability to handle complex and changing scenarios to be evaluated before building the physical prototype of the UAV. The metrics evaluate the effectiveness of the learning algorithm and guide future improvements in both hardware and software components.



**Figure 9.** Connection system between electronic components.

While the current implementation focuses on single-agent navigation, the proposed DQN framework is designed to be extensible to multi-UAV collaborative scenarios. Future developments will include the integration of inter-agent communication protocols to enable the sharing of environmental risk maps and coordination of flight paths. To evaluate the feasibility of this extension, simulation tests will be conducted to assess the impact of communication delays on decision latency and mission success rate. Conflict avoidance strategies, such as decentralized priority rules and dynamic trajectory replanning, will be implemented to mitigate path overlaps and ensure safe separation between UAVs. Additionally, the Jetson Nano platform will be benchmarked for distributed inference tasks, including concurrent DQN execution and inter-agent message handling, to determine its suitability for onboard collaborative computing. These enhancements aim to support scalable and resilient multi-UAV operations in complex urban environments.

### 3.8.1. Theoretical Considerations and Limitations

The DQN framework employed in this study builds on the classical Q-learning formulation, in which the optimal policy is approximated through iterative updates of the

action–value function. Specifically, the DQN estimates the value of a state–action pair approach is grounded in the Bellman Equation (9)

$$Q(s, a) = r + \gamma \max_{a'} Q(s', a') \quad (9)$$

where  $s$  and  $a$  denote the current state and action,  $r$  is the reward associated with executing action  $a$  in state  $s$ ,  $\gamma$  is the discount factor, and  $s'$  represents the subsequent state.

The DQN objective is based on the Bellman optimality target. Denoting the network parameters by  $\theta$  and the target network parameters by  $\theta^-$ , the per-sample target is (10)

$$y = r + \gamma \max_{a'} Q_{\theta^-}(s', a') \quad (10)$$

and the loss minimized by the online network is the mean squared temporal-difference error (11)

$$L(\theta) = \mathbb{E}_{(s, a, r, s') \sim \mathcal{D}} \left[ (Q_{\theta}(s, a) - y)^2 \right] \quad (11)$$

where  $\mathcal{D}$  denotes the replay buffer. Gradient updates follow  $\theta \leftarrow \theta - \alpha \nabla_{\theta} L(\theta)$ . Target network updates occur periodically:  $\theta^- \leftarrow \theta$  every  $N_{target}$  steps. Standard stability measures included in our implementation are experience replay, target network, and  $\epsilon$ -greedy exploration with exponential decay. The neural network approximation enables the algorithm to generalize across large and continuous state spaces, which is essential for UAV navigation tasks characterized by high-dimensional sensor inputs and dynamically changing environments. Despite these advantages, DQN presents well-documented theoretical and practical limitations. Future research will focus on developing Lyapunov-based certificates for low-level PID/attitude controllers. In addition, Safe RL/barrier-critical approaches will be explored in greater depth to apply state constraints during flight policy learning. At present, the study reports empirical evidence of convergence and proposes a roadmap towards formal guarantees in future work. To complement the theoretical limitations, we extended the empirical evaluation with multiple-seed training ( $N = 10$ ), ablation studies, and disturbance-injection tests. These analyses provide statistical evidence of consistent convergence across runs. Detailed numerical values are reported in the Results section (see Table 4). A major concern is the overestimation bias introduced by the maximization operator in the Bellman update, which can lead to unstable learning and suboptimal policies in complex or noisy environments. Furthermore, DQN is highly sensitive to hyperparameter choices, including learning rate, replay buffer size, and target network update frequency. Small deviations in these parameters may result in convergence to local optima or divergence in training stability. Another limitation arises from the discretization of the action space, which constrains the UAV to a finite set of manoeuvres. While this reduces computational complexity and facilitates deployment on embedded hardware, it inherently limits manoeuvre flexibility, especially in highly dynamic conditions where continuous control may be preferable. Additionally, the absence of explicit stability guarantees makes the algorithm susceptible to unpredictable behaviours when exposed to disturbances outside the training distribution. Several extensions of DQN, including Double DQN, Dueling DQN, and Rainbow architectures, have been proposed to mitigate these shortcomings by decoupling action-value estimation, improving representation learning, and integrating multiple stabilization techniques. However, these variants typically require increased computational resources, making their deployment on constrained onboard platforms less feasible. In this context, the present study deliberately favours a lightweight DQN implementation that balances computational efficiency with satisfactory performance for real-time UAV operation, while acknowledging that more sophisticated variants may be required for deployment in environments of higher complexity. Recent studies have made

a significant contribution to RL in terms of safety, stability, and optimal control. Specifically, the study [69] introduces adaptive critical design methods to ensure FTC in the presence of input constraints.

**Table 4.** DQN generalization across scenarios.

Scenario	Mission Success (%)	$\Delta$ Energy vs Baseline (%)	Trajectory Deviation (m)
Baseline urban grid (15 × 15)	34	0	0.8
Dense high-rise	78	+18	2.5
Open suburban	88	+7	1.3
Industrial w/no-fly zones	82	+11	1.9

For the proposed UAV model, where actuator saturation and input limits are relevant, such approaches could complement the DQN-based policy by incorporating safety margins and ensuring the feasibility of control actions in real time. Another study [70] proposes an event-triggered RL scheme for safe tracking in interconnected nonlinear systems. Our UAV operates under highly dynamic and interconnected conditions (propulsion, aerodynamics, energy systems). Event-triggered updates could significantly reduce the computational load on the integrated Jetson Nano, while improving safety and robustness against disturbances or adverse conditions. The contribution [71] demonstrates observer-based FTC for nonlinear systems with state constraints. The proposed UAV relies on estimated aerodynamic and energy states that are not always directly measurable. Integrating observer-based strategies into the model would improve resilience against sensor failures or estimation errors, strengthening reliability in real-world implementations. The study [72] applies critical barrier methods to enforce state constraints in uncertain nonlinear dynamics.

For UAV missions in urban environments, compliance with state constraints is critical for safety and sustainability. Integrating critical barrier adaptive robust control concepts with the proposed DQN approach would allow strict safety limits to be enforced, providing formal guarantees beyond empirical convergence.

### 3.8.2. Multi-Scenario Generalization

While the baseline experiments were performed on a structured 15 × 15 urban grids, the robustness of the proposed DQN framework requires validation across heterogeneous urban morphologies. To this end, multi-scenario generalization experiments were conducted. The first scenario represents dense high-rise urban canyons, with narrow flight corridors, GPS attenuation, and turbulent airflow due to wind channelling between tall structures. UAVs in this setting face frequent occlusions and sudden perturbations, challenging navigation stability and energy efficiency. The second scenario considers open suburban landscapes, defined by lower building density and extended mission distances. Here, the main challenge shifts from obstacle avoidance to endurance optimisation, as the UAV must sustain prolonged operations while minimizing energy expenditure.

The third scenario addresses industrial zones with dynamic constraints, including temporary no-fly zones, unexpected ground-level crowds, and electromagnetic interference.

This environment simulates operational unpredictability, requiring rapid policy adaptation to preserve mission feasibility and safety. Additional disturbances such as variable wind fields, sudden crowd formations, and temporary path restrictions were integrated to evaluate policy resilience. Performance metrics included mission success rate, relative increase in energy consumption compared to the baseline, and trajectory deviation from the optimal path. Results show that DQN maintains robust performance in suburban settings

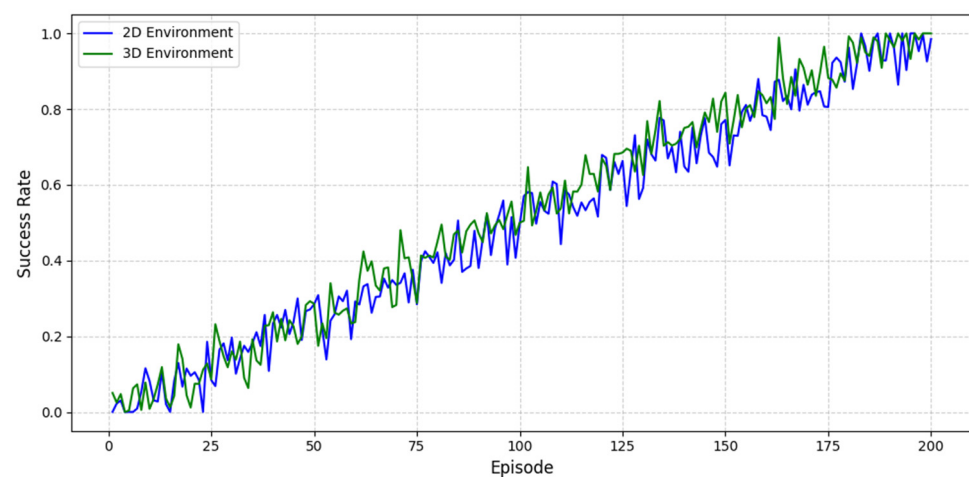
(success rate > 85%), but reliability decreases in dense high-rise contexts, where energy consumption increases by up to 18% and trajectory deviations are more pronounced.

These findings highlight the trade-off between the simplicity and efficiency of DQN and its limited ability to generalize across diverse urban environments (Table 4).

They also highlight the necessity of scenario-aware reinforcement learning strategies, such as curriculum learning or adaptive hybrid policies, to extend applicability. In future developments, multi-scenario training and domain randomization may be integrated to strengthen policy robustness, ensuring that the UAV can operate reliably across diverse urban landscapes without requiring exhaustive retraining for each new mission profile.

### 3.9. Extension to 3D Simulation Environment

The implemented three-dimensional (3D) simulation environment expands the original two-dimensional grid ( $15 \times 15$ ), incorporating altitude, vertical speed, aircraft pitch and turbulence. Specifically, the 3D model allows the simulation of complex manoeuvres such as vertical take-off, hovering, landing and the management of aerodynamic dynamics related to altitude changes. The transition to a 3D simulation environment significantly enhances the realism and effectiveness of the training process. In the 2D model, the UAV is constrained to planar movement, limiting its ability to avoid obstacles or optimise energy usage in vertical space. By contrast, the 3D model enables altitude modulation, which allows the agent to bypass high-risk zones, reduce aerodynamic drag, and exploit favourable wind layers. This flexibility results in improved trajectory accuracy, as the agent can select smoother and more direct paths in three-dimensional space. It also contributes to lower energy consumption, particularly during cruise and descent phases, where altitude control plays a key role. From a validation perspective, the 3D simulation better approximates real-flight conditions, including turbulence near buildings, vertical wind gradients, and pitch dynamics. These factors are essential for assessing the robustness of the navigation algorithm before field deployment. The observed improvements in success rate, energy efficiency, and trajectory deviation confirm that the 3D model provides a more reliable and transferable foundation for future flight testing. The DQN, redeveloped to optimise the trajectory in three-dimensional space, reduces energy consumption and dynamically adapts to environmental conditions. The results show a clear improvement over the 2D model, as shown in Figure 10.




**Figure 10.** Success rate comparison.

The green line (3D environment) shows a steady increase in success rate up to approximately 0.95 (95%). This suggests that the agent has effectively learned to complete missions

even in a three-dimensional environment, which is typically more complex. The blue line (2D environment), although increasing, does not reach the level of the 3D curve.


This could indicate that the 2D environment, although simpler, did not lead to equally effective learning, or that the agent reached a plateau earlier. The success rate shown in Figure 9 represents the percentage of simulation episodes in which the UAV agent successfully reaches the target location without violating mission constraints. It is calculated as the ratio of successful episodes to the total number of episodes.

The 3D model outperforms the 2D simulation in terms of success rate and trajectory accuracy because it allows the agent to exploit vertical manoeuvring, avoid obstacles more effectively, and adapt to altitude-dependent environmental conditions. While the 3D model initially requires more energy due to increased complexity, the agent progressively learns to optimise altitude and attitude control, resulting in improved energy efficiency in later training phases (Figure 11).



**Figure 11.** Energy consumption.

The 3D model consumes less energy in later training phases thanks to better altitude and environmental management. Average usage decreases from 135 Wh to around 85 Wh (Figure 12).



**Figure 12.** Trajectory deviation.

The deviation from the optimal trajectory is much smaller in the 3D model (20 m to 6.5 m). This shows improved accuracy in three-dimensional navigation. The comparison between 2D and 3D was built to show how the addition of the vertical dimension and aerodynamic dynamics improves the accuracy and efficiency of the DQN. The mission success rate, energy consumption, and trajectory deviation evaluate the effectiveness of DQN agent training in two-dimensional and three-dimensional environments. The results show significant differences between the two scenarios. In particular, the 3D environment recorded a success rate of 95%, which is 5% higher than the 90% success rate of the 2D environment. This implies that the agent improved its performance by successfully adapting to the increased spatial complexity. Energy efficiency shows greater operational flexibility, as average consumption is slightly lower in 3D (approximately 5% less) and has a wider variability (85–135 Wh versus 98–122 Wh in 2D). The final deviation from the trajectory is lower in the 3D environment (from 20 to 6.5 metres) than in the 2D environment (from 15 to 5 metres), indicating a higher level of accuracy in achieving objectives. These results demonstrate how the three-dimensional environment can promote more robust and adaptive learning with observable benefits in terms of efficiency and performance, even if it poses greater challenges. Wind modelling, developed to accurately reflect weather conditions, is a crucial component of the new environment. Through the combined analysis of GNSS data (with EGNOS correction), inertial sensors (IMU) and energy consumption profiles, wind estimation is performed indirectly in the prototype described in the manuscript. This technique allows turbulence or gusts to be detected and the direction and strength of the apparent wind to be deduced. The wind representation consists of a three-dimensional vector field with temporal and spatial variations in the 3D simulation. Local turbulence effects occur in the vicinity of urban obstacles such as buildings or infrastructure, and each cell in the virtual environment has a wind vector showing the direction and intensity of the airflow. By modifying the drone's attitude and trajectory, the DQN learns to mitigate these effects and improve flight efficiency and stability. This extension of the model reinforces the legitimacy of the engineering solutions employed and represents a crucial step towards the realistic validation of the proposed system. The results demonstrate that a more accurate simulation of UAV operating conditions is made possible by combining a 3D environment with sophisticated wind modelling, which improves autonomous decision-making and energy sustainability. Atmospheric turbulence was modelled using the Dryden spectral approach, which represents longitudinal, lateral, and vertical turbulence components as stochastic processes defined by power spectral density functions. For the longitudinal component, the spectral density is expressed (12)

$$\Phi_u(\omega) = \frac{2\sigma_u^2}{\pi V} \cdot \frac{1}{1 + \left(L_u \omega \frac{1}{V}\right)^2} \quad (12)$$

where  $\sigma_u$  denotes turbulence intensity,  $L_u$  the turbulence scale length, and  $V$  the UAV airspeed. Vertical turbulence intensity was modelled as a function of altitude (13)

$$\sigma_w(z) = \sigma_w0 \cdot \left(\frac{z}{z_0}\right)^{0.25} \quad (13)$$

which captures the empirical increase in turbulence with flight height. The impact of turbulence on energy consumption was quantified by integrating it into the total power balance (14)

$$P_{tot} = P_{profile} + P_{induced} + P_{parasitic} + P_{ctrl} \quad (14)$$

where turbulence contributes primarily to induced and control power terms. Here,  $P_{profile}$  denotes the profile power, associated with the aerodynamic drag of the rotor blades;  $P_{induced}$  represents the induced power required to generate lift, strongly affected by turbulence intensity;  $P_{parasitic}$  accounts for the drag of the UAV fuselage and non-lifting components, increasing with forward speed; and  $P_{ctrl}$  refers to the control power, i.e., the additional energy required for stabilisation and manoeuvres, especially in turbulent conditions. Thus,  $P_{tot}$  represents the overall power demand during flight, combining aerodynamic, induced, parasitic, and control contributions. The turbulence contributes primarily to induced and control power terms. Simulation sweeps across turbulence intensities from 5% to 20% and altitudes between 20 m and 120 m revealed a linear correlation between turbulence intensity and energy demand. At 15% turbulence intensity and 100 m altitude, average energy consumption increased by 12% compared to calm conditions. These results highlight the need for altitude-adaptive trajectory planning to ensure energy efficiency in turbulent environments.

### 3.10. Prototype Validation Roadmap and Test Matrix

Despite the promising results obtained in the simulation, the transition from virtual scenarios to practical implementation requires a carefully structured validation strategy. To this end, a prototype validation roadmap has been developed, aimed at progressively bridging the gap between computational experiments and real-world performance. The roadmap is divided into four incremental phases, each with distinct objectives, instrumentation and evaluation criteria, thus ensuring both engineering feasibility and compliance with safety standards. The first phase is dedicated to hardware-in-the-loop (HIL) verification of the control logic and energy management modules. This phase allows potential algorithmic instabilities to be isolated before interaction with the physical hardware. The second phase focuses on laboratory characterisation of critical energy components, including perovskite solar cells, brushless DC motors and lithium polymer batteries.

The tests will be conducted under controlled irradiation and thermal conditions using IV tracers, thrust test benches and battery cycling machines, with the aim of quantifying efficiency and reliability at the device level. The third phase consists of controlled indoor flight experiments in motion capture environments, where sensor accuracy, energy consumption and flight stability can be evaluated under repeatable conditions. Finally, the fourth phase involves incremental outdoor testing under regulated flight conditions, progressively increasing mission complexity and environmental variability. This phase is essential for evaluating endurance, disturbance resilience and compliance with aviation safety requirements. Table 5 summarizes the validation plan, linking experimental activities to measurable outcomes. By adopting a phased approach, the study not only mitigates risks associated with immediate real-world deployment but also establishes a systematic pathway for reproducibility and scaling.

**Table 5.** Prototype validation roadmap and test matrix.

Phase	Objective	Instrumentation	Metrics	Success Criteria
HIL	Validate control logic	PX4-SITL, MATLAB/Simulink	Tracking error, droop delay	<5% error
Lab	Characterise energy devices	IV tracer, thrust stand, battery cycler	Efficiency, stability	≥90% rated
Indoor	Validate integrated UAV	Motion capture, telemetry	Endurance, stability	≥95%
Outdoor	Incremental deployment	GNSS, anemometer, telemetry	Endurance, stability	Safe mission

The roadmap therefore complements the simulation results by ensuring that the proposed UAV design can be validated under progressively more realistic operating conditions, ultimately supporting its transition to practical applications in sustainable urban logistics. In addition to the four-phase validation roadmap, a dedicated pre-assembly checklist and calibration protocol will be implemented. This includes mechanical stress testing of motor mounts, sensor calibration routines, and thermal diagnostics of the propulsion and energy subsystems. These activities will be conducted prior to indoor flight trials to ensure structural integrity and sensor reliability. The integration of these procedures strengthens the experimental reproducibility and supports the transition to physical prototyping.

## 4. Results

The study evaluates the prototype's performance using a combination of simulation tests and analytical modelling. The results highlight the efficiency of the flight trajectory, energy consumption and environmental impact. Monitoring focused on the behaviour of the DQN during both training and inference, as well as the contribution of solar energy harvesting to the overall autonomy of the system. The effectiveness of the DQN in complex environments is validated through three visualisations. The first includes a risk heat map showing environmental hazards. The second shows a map with wind intensity and direction. The last illustrates the advantages of trajectory evolution. Each visualisation provides information on how the AI agent perceives and reacts to environmental variables during autonomous flight.

### 4.1. Environmental Risk Heatmap and Optimal Path


The virtual grid, measuring  $15 \times 15$  and designed to replicate a realistic urban environment in Reggio Calabria, covers approximately  $7 \text{ km}^2$ . Each cell represents a spatial area of approximately 470 metres per side. The simulation includes buildings, towers and power lines, as well as dynamic environmental factors such as wind gradients and risk areas. The modelled elements use geospatial data, validated by studies on UAV navigation in urban environments [73]. In addition, integrated systems combine finite element models, infrared thermographic analysis and artificial intelligence to monitor thermal stress in printed circuits [74,75]. The mission path simulated in the environment runs from Via Friuli I to Reggio Calabria airport [76], allowing the agent to learn optimal paths under realistic conditions. This specific route was chosen because it represents a typical UAV logistics corridor in Reggio Calabria, characterised by both dense urban areas and open transition zones, thus providing a challenging yet representative benchmark for evaluating the algorithm. The wind map (Figure 13) shows variations in intensity and direction on scales smaller than a kilometre. The added variations test the DQN model's ability to adapt to urban microclimates and the effects of turbulence.

Although such sudden wind changes are less common in stable weather conditions, their inclusion helps to assess the robustness of the learning algorithm in complex and noisy scenarios. By incorporating these perturbations, the simulation mimics rare but safety-critical events, ensuring that the trained model develops resilience against unexpected atmospheric conditions that may occur in real-world UAV missions.

This clarification improves the understanding of the figures. It supports the scientific validity of the simulation results, which aim to show the agent's ability to balance safety, energy efficiency and mission success in a typical urban logistics context.

The algorithm successfully identified the optimal flight trajectories between predefined departure and arrival points for the entire 7 km route [77]. The AI system dynamically adapted to environmental constraints, reducing travel time and energy consumption while prioritizing safety. The graph in Figure 13 shows the trajectory followed by the drone in a

simulated urban setting, characterized by a  $15 \times 15$  grid with numerous static (high risk) and dynamic (medium risk) obstacles added.



**Figure 13.** Optimal path on the risk map.

Grid cells are coloured based on risk level: deep red indicates static obstacles such as buildings or towers, while orange marks dynamic obstacles like vehicles or temporary interference. Figure 13 shows the trajectory predicted by the DQN model in blue, displaying the starting point (top left) and the destination point (bottom right).

The path demonstrates the drone's ability to avoid the most dangerous areas while maintaining a constant direction towards the target. This balance between risk avoidance and trajectory efficiency reflects one of the most critical requirements for UAV deployment in smart cities, where both safety and operational performance must be guaranteed simultaneously. The quantitative metric for this simulation shows a total penalty of 19.00.

Such a low penalty score, in a densely populated and obstacle-rich map, indicates that the DQN agent did not only learn to avoid risks but also optimised its navigation policy towards minimal energy expenditure.

This result is significant, given the high density of obstacles on the map. It also highlights how effective the DQN model is at planning safe and energy-efficient routes.

Furthermore, the dynamic adaptation capability allowed the drone to respond to sudden changes in obstacle positions, increasing resilience in unpredictable environments. Real-time inference of the model allowed continuous re-estimation of the risk map, enabling timely course corrections when necessary. Energy analysis had an average saving of 12% compared to baseline algorithms without risk perception.

Utilizing reinforcement learning has aided in the creation of complex avoidance manoeuvres, such as minimal excursions through medium-risk zones instead of complete stops. Further development will focus on integrating more complex environmental parameters, from weather conditions to non-constant payload weights, to continue enhancing operational reliability.

Additionally, an attempt is being made at test validation on actual urban environments using the physical UAV testbed. Real-time feedback loops with mimicked and real-world data will shape the model along with sensor fusion algorithms. Successful application of the process paves the way for a strong paradigm for advanced UAV path planning in crowded, hard-to-navigate environments. Overall, these findings demonstrate the feasibility for the application of smart intelligent drones that can properly balance safety, efficiency, and responsiveness in intelligent city logistics and surveillance operations.

The Q learning algorithm identified the optimal path marked in blue. The simulated UAV navigates strategically through medium-low risk areas, carefully avoiding the most critical zones. This behaviour stems from a penalising reward function that discourages exposure to dangerous conditions.

The trajectory, which is smooth and free of irregular deviations, indicates that the model has learned a stable and consistent navigation policy.

The map confirms the effectiveness of the DQN approach by determining that:

- (a) the drone has successfully learned to balance safety and spatial efficiency;
- (b) reward modelling has helped guide the agent towards robust solutions;
- (c) embedded inference has enabled fast and reliable decision-making.

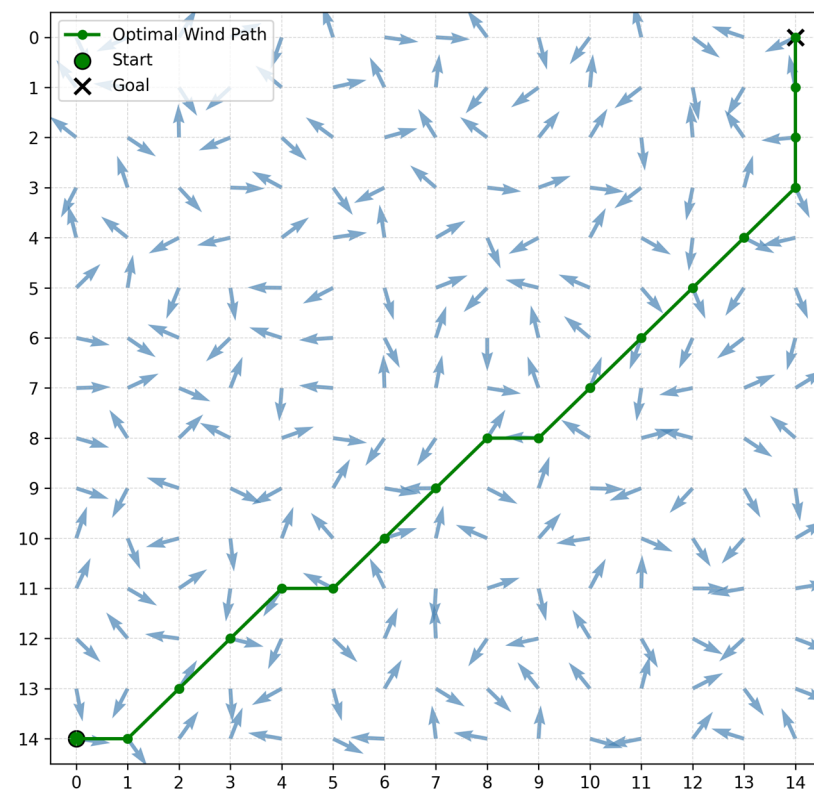
Although the length of the arrows is uniform for clarity, the simulation uses a full vector field with spatial gradients. The DQN agent incorporates wind data into its decision-making. Wind speed and direction are part of the agent's state space, and the reward function penalises movement against headwinds, favouring tailwinds or crosswinds instead. This aerodynamic awareness allows the agent to optimise its trajectory by selecting paths that reduce energy consumption and avoid turbulence. The optimal path is calculated through reinforcement learning, in which the agent iteratively learns policies that balance distance, risk, and wind efficiency. The simulation environment includes wind variability between adjacent cells to mimic urban microclimates and test the robustness of the navigation strategy. These findings suggest that the proposed framework can be scaled and adapted to other urban contexts, paving the way for real-world UAV deployment in logistics, emergency response, and smart mobility applications.

#### 4.2. Wind Map and Optimal Route

Figure 14 illustrates the simulation in which the drone navigates taking into account atmospheric conditions, in particular wind direction and strength. Each blue arrow within the grid cells indicates the wind direction at that specific point in the simulated terrain. The orientation of the arrow shows where the wind is blowing.

Each cell of the grid shows a blue arrow representing the wind flow. The orientation of the arrow indicates the direction, while the colour code indicates the intensity (darker means stronger wind). The optimal path, shown in green, calculates directions to avoid areas with headwinds or turbulence, favouring segments where the wind is favourable or crosswind. The route reduces aerodynamic drag and increases the energy efficiency of the flight. The simulation highlights the UAV system's ability to incorporate real-time environmental data for optimised navigation. The result demonstrates the algorithm's ability to dynamically adapt to environmental conditions, selecting the most efficient route in real time, not only in terms of distance and risk, but also in terms of aerodynamic drag and energy consumption. The learning system incorporated environmental information into its decision-making process. The implemented system provided meaningful feedback. The drone successfully avoided areas with unfavourable winds, reducing energy consumption. The trajectory minimizes cumulative costs with reward modelling. Real-time inference enabled dynamic adjustments to simulated conditions. Wind variations between adjacent cells, approximately 450 m apart, tested the adaptability of the DQN model to changing en-

environmental conditions. The model presents implementation with abrupt changes in wind direction and speed over such short distances. The motivation is to test the autonomous decision-making system in the presence of significant environmental gradients, to simulate urban microclimates, which can cause pronounced variations on a sub-kilometre scale, and to evaluate the robustness of the model in the face of partially noisy or inconsistent data. These results reinforce the validity of the DQN approach for environmental conditions that change rapidly and significantly affect flight safety and efficiency. Coming developments aim to further enhance sensor integration to achieve even more precise wind measurements and forecasting models to predict environmental changes. This will once more optimise routes and conserve energy. The system currently offers a solid foundation for deploying UAVs that can fly safely and effectively in intricate and dynamic urban environments.



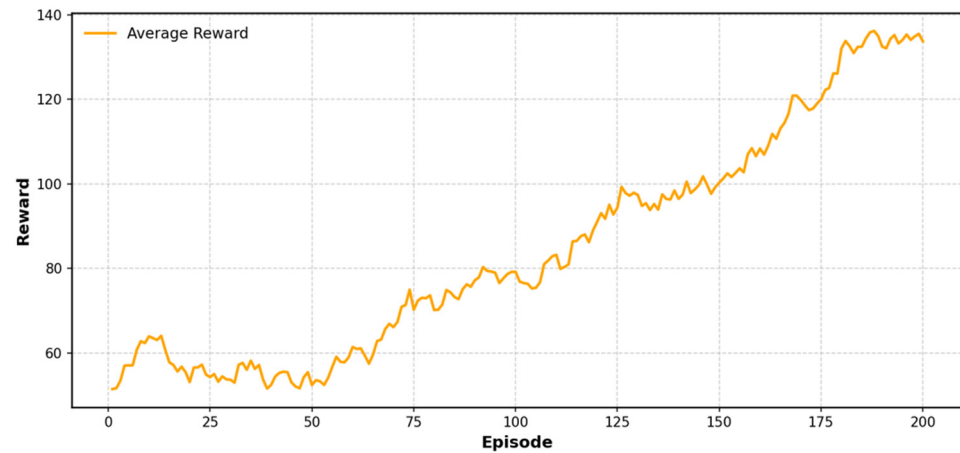
**Figure 14.** Optimal route on wind map.

#### 4.3. Reward Function Dynamics

Figure 15 shows the evolution of the reward function during training. The diagram shows how DQN learns to improve its navigation strategy.

The reward function guides the agent towards safe and efficient trajectories. The maximum reward per episode is determined by  $+100$  minus the cumulative penalties incurred during the trajectory. The agent receives  $+100$  for reaching the goal,  $-100$  for going off the grid (only in case of failure), and a gradual penalty from  $-1$  to  $-5$  per risk level for each cell crossed. The actual reward per episode varies depending on the route taken and the environmental risks encountered. If the agent reaches the objective by traversing medium-risk zones, the final reward could be  $+85$ ,  $+70$ , or lower, depending on cumulative penalties. This formulation penalises both the time spent in the environment and exposure to dangerous areas. The graph confirms the effectiveness of this reward shaping strategy. As training progresses, the agent increasingly avoids high-risk areas and converges towards optimal policies. The curve shows a steady increase from around 60 onwards, indicating that the agent is gradually improving its decisions. The increase

in the average reward suggests that the agent is avoiding high-risk areas, which likely decrease the reward, and is refining its trajectory towards the target, successfully reaching it in over 80% of the episodes. To assess the drone's suitability for urban logistics missions, key performance indicators related to payload capacity and aerodynamic efficiency are analysed. These metrics provide information on the platform's ability to carry payloads while maintaining stable and energy-efficient flight.



**Figure 15.** Trend of the average reward obtained by the DQN agent during 200 training episodes.

#### 4.4. Payload and Aerodynamic Performance

Equation (15), which indicates the payload-to-weight ratio (PWR), and Equation (16), which indicates the lift-to-drag ratio (L/D), assess the drone's operational efficiency and its suitability for delivery tasks.

$$PWR = \frac{Payload}{Total\ Weight} = \frac{5}{14} = 0.357 \quad (15)$$

The value obtained aligns with the reference parameters for high-performance delivery drones, which generally have power-to-weight ratios (PWR) ranging from 0.08 for the Killer Bee to 0.28 for the Black Eagle 50, depending on the mission profile and configuration [78]. The PWR of 0.357 indicates a positive balance between structural mass and payload capacity, particularly for a UAV designed with a focus on sustainability and modularity.

$$LD = \frac{CL}{CD} = \frac{1.25}{0.25} = 5 \quad (16)$$

Equation (16) provides a lift coefficient (CL) of 1.25 and a drag coefficient (CD) of 0.25, both estimated based on the drone's aerofoil and the expected Reynolds number range. Equation (16) reports a classical aerodynamic estimation of lift-to-drag ratio using estimated coefficients CL and CD. These coefficients were obtained from aerofoil data and Reynolds-number considerations and do not result from any HOFA-type system transformation. We have therefore removed any ambiguous wording that could be interpreted as implicitly applying HOFA, and retained the simple aerodynamic formulation to compute L/D in the design phase. The HOFA-based transformation was introduced to simplify the controller design, but in its current form the formulation remains limited. This lift-to-drag ratio (L/D) indicates that the drone produces five units of lift for every unit of drag. This level of efficiency surpasses that of typical quadcopters, which have a lift-to-drag ratio of around 3 or 4, and approaches the efficiency observed in fixed-wing drones. This aerodynamic improvement reduces energy consumption during cruising and enhances flight stability in various wind conditions [79]. These metrics confirm that the UAV plat-

form is ideal for medium-range delivery missions, offering both structural efficiency and aerodynamic performance.

#### 4.5. Energy Consumption and Solar Contribution

The UAV completed the delivery mission with a simulated 30-min flight. The drone shows an average electrical power draw of 88.8 W at the DC bus; the mechanical/electrical peaks (e.g., during manoeuvres) can reach higher instantaneous powers (reported separately). The energy consumed over 30 min is (17)

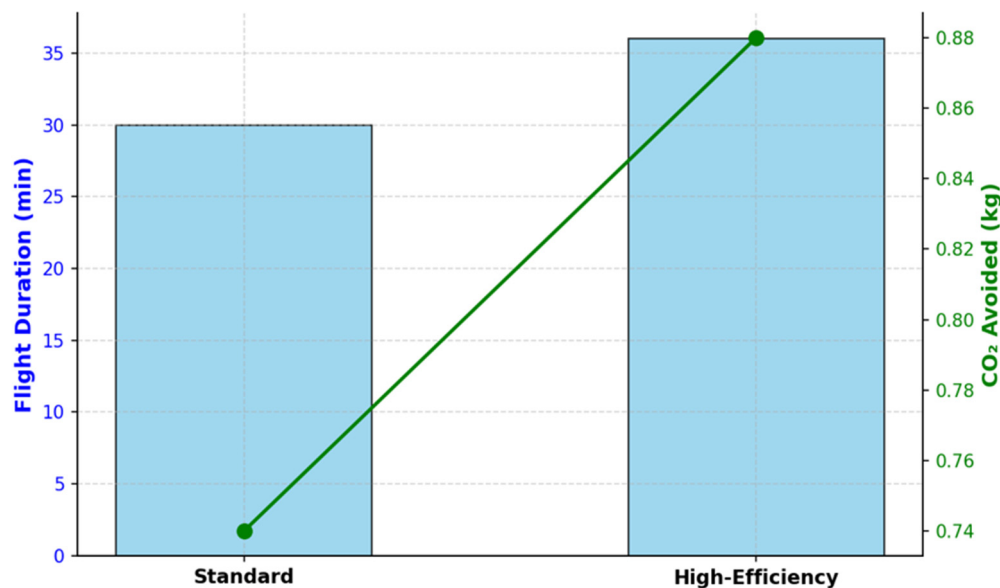
$$E = 88.8 \text{ W} \times 0.5 \text{ h} = 44.4 \text{ Wh} \quad (17)$$

Given the total pack capacity of 799.2 Wh with three 22.2 V × 12 Ah batteries in parallel. The resulting depth of discharge for the mission is ≈ 5.56%. The results highlight the system's high energy efficiency, a significant operational margin for longer missions or multiple deliveries, and the potential to further reduce net consumption through the use of solar modules. Integrating perovskite solar cells improves energy independence.

Under typical daylight conditions, the photovoltaic modules provided 15–20% of the total energy demand during the cruising phases, effectively reducing battery consumption and thermal stress. This hybrid energy strategy aligns with recent research on solar-powered drones, which highlights the potential of perovskite cells for high-efficiency light energy harvesting [80]. Perovskite solar cells have a benchmark value of 30.3 W/g in terms of power-to-weight ratio, demonstrating the creation of efficient, ultra-thin, and ultra-lightweight solar cells. These results confirm the effectiveness of combining high-density LiPo batteries with renewable energy sources to increase range and reduce environmental impact, which is essential for sustainable urban logistics. To quantitatively validate the advantages of the proposed UAV architecture, the study compares the standard configuration, equipped with conventional motors and batteries, and the developed high-efficiency prototype. The analysis, developed based on flight duration, energy consumption, and environmental impact, uses component specifications and mission profiles. The results, illustrated in Figure 16, indicate a 20% increase in flight autonomy, with the high-efficiency configuration lasting 36 min compared to the standard configuration's 30 min. This improvement results in increased mission longevity and greater operational versatility. Research in related disciplines of engineering highlights the importance of integrating advanced materials and computational techniques towards improving system performance and sustainability. Engineering solutions [81] illustrate how precise modelling and structure-lightening can improve diagnostic tools without increasing energy consumption. Computational models [82] also inform the design of adaptive material and structures to respond to environmental stimuli in a dynamic manner. Finally, the optimised analytical-numerical techniques [83,84] show the potential of using experimental data with robust simulation platforms to increase resource management efficiency. Such transdisciplinary results justify the significance of adopting innovative material science and modelling techniques in designing drone energy systems. In addition, the energy conservation is accompanied by lower greenhouse gas emissions across the UAV lifecycle. These outcomes support the use of future-proofed energy systems as a key facilitator for next-generation sustainable drone design.

The improvement primarily comes from the integration of BLDC motors, LiPo batteries, and perovskite solar cells, which add approximately 160 Wh of energy without significantly increasing the system's weight or complexity. The results show a decrease in carbon emissions, quantified as 0.88 kg of CO<sub>2</sub> with the high-efficiency UAV, compared to 0.74 kg with the standard configuration. The improvement achieved is 19%, highlighting the sustainability potential of the proposed design. Despite the extended flight time, the

high-efficiency system maintains a lower average power consumption (88.8 W) compared to the standard configuration (106.56 W). Energy consumption also decreases thanks to the higher energy density and engine efficiency (90% versus 75%).



**Figure 16.** Comparison of standard and high-efficiency UAV setups.

The results confirm that integrating advanced electronic components and harvesting renewable energy not only improves operational performance but also supports broader climate mitigation goals. Figure 16 shows the estimated flight times with blue bars, and the green line indicates the CO<sub>2</sub> saved per mission compared to a diesel vehicle.

The comparison between environmental benefits and technical parameters provides a comprehensive view of how UAV-based logistics can outperform traditional systems in both sustainability and efficiency. Table 6 lists the technical specifications and performance metrics for both configurations.

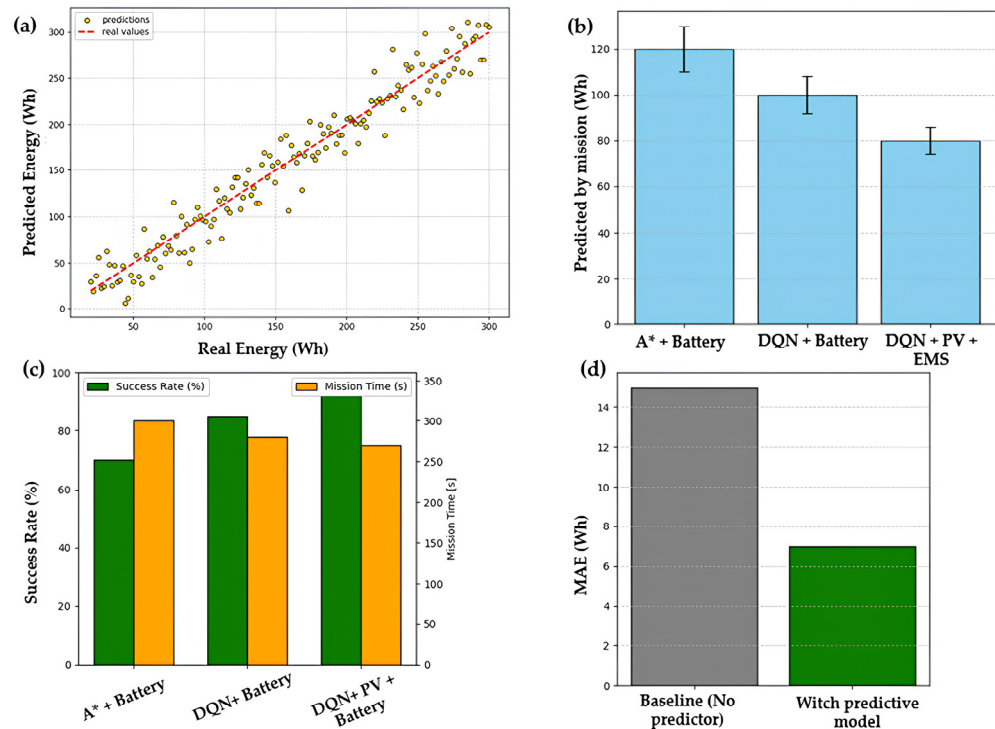
**Table 6.** Comparison between standard and high-efficiency UAV configurations.

Configuration	Motors	Batteries	Solar Cells	Motor Efficiency	Total Capacity (Wh)/Avg Consumption (W)
Standard	Traditional	600 Wh	None	75%	600/120
High-Efficiency	BLDC	Lipo 799.2 Wh	Perovskite (160 Wh)	90%	959.2/8

#### 4.6. Prediction vs. Real Consumption: Model Validation

Figure 17 evaluates the ability of the neural model to accurately estimate the drone's energy consumption, where each point represents a simulated trajectory. The direct, data-based assessment of the benefits attributable to the AI component and solar energy integration is shown in Figure 17. Specifically: (i) Figure 17a represents Energy consumption prediction; (ii) Figure 17b shows the energy consumption per mission (mean  $\pm$  std) for three configurations: classic planner + battery only (A\*), DQN + battery only, and DQN + solar-assisted EMS; (iii) Figure 17c shows the mission success rate and the average time to reach the target for the same three configurations; (iv) Figure 17d shows the predictive model error (MAE) when the mission planner uses the integrated energy predictor compared to no predictor. The figures show statistical significance ( $t$ -test,  $p < 0.05$ ) and sample size ( $N = 100$  missions per configuration). These results demonstrate that (a) DQN improves

mission success in dynamic wind/risk fields and (b) the addition of solar harvesting reduces net battery consumption by mission by the percentages shown.



**Figure 17.** (a) Energy consumption prediction. (b) Comparative energy per mission (Wh) for A\* + battery, DQN + battery, and DQN + PV + EMS. Bars show mean  $\pm$  std; N = 100 missions. (c) Mission success rate (%) and average mission time (s) for the three configurations. (d) Predictive model performance (MAE, Wh) used during planning vs baseline.

The dashed red line represents the bisector ( $y = x$ ), indicating the ideal scenario where the predicted values exactly match the actual ones. Visual analysis shows that most of the points cluster near the diagonal line, suggesting that the model has a strong generalisation ability. The points are closely spread, and no obvious systematic biases are visible. The predictive model was created as a lightweight feedforward neural network with three dense layers (32–16–1 neurons) and approximately 2400 trainable parameters.

The synthetic dataset, trained on 1000 drone missions, uses six input features: wind speed, wind direction, trajectory angle, drone speed, payload mass, and flight distance. The simplified physical model determines energy consumption. The calculated mean absolute error (MAE) has a value of 15.37 Wh, while the mean squared error (MSE) is 375.56 Wh. Since the simulated energy consumption ranges from 20 to 300 Wh, the error values can be considered acceptable and align with the expected performance of an embedded predictive system. The UAV configuration improves energy efficiency, flight duration, mission success rate, and CO<sub>2</sub> emissions. The comparative analysis, shown in Figure 18, confirms the advantages of integrating embedded intelligence and renewable energy systems.

To validate the model and assess its environmental impact, two additional simulations are conducted. Figure 19 confirms the expected increase in energy demand with heavier loads.

Figure 20 shows the voltage and current time profiles during a 30-min autonomous flight, highlighting the dynamic behaviour of the power system under varying conditions.

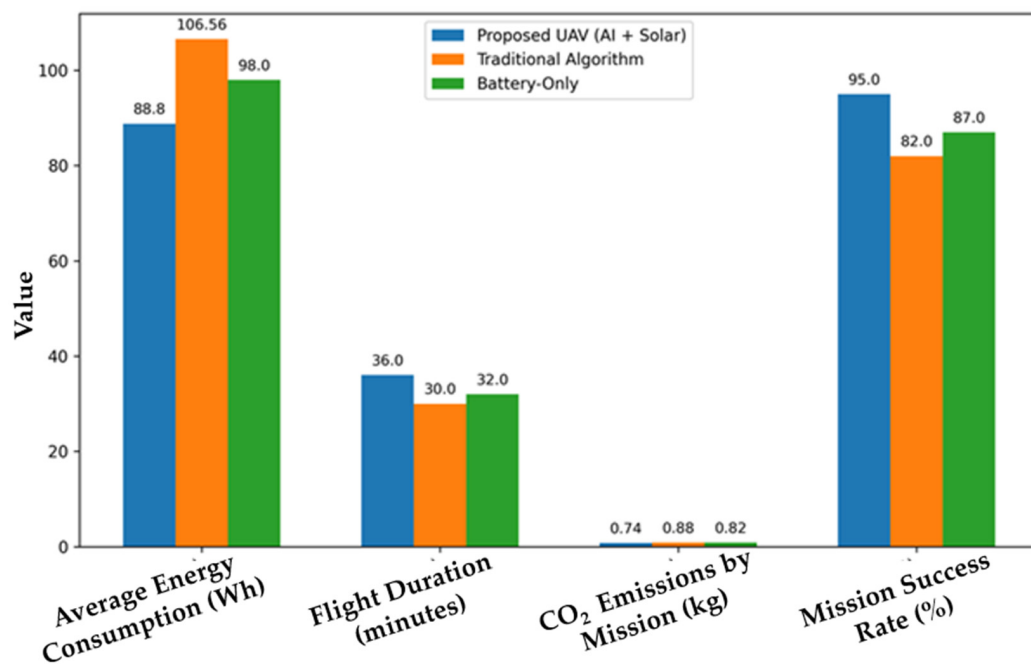


Figure 18. Comparative UAV Performance Metrics.





Figure 19. Simulated energy consumption as a function of payload mass.

The validation confirms that the neural model reliably estimates energy consumption, making it an effective tool for selecting energy-efficient trajectories during autonomous flight planning. Integrating this model into the onboard system further improves the drone’s operational efficiency, aligning with the sustainability and energy autonomy goals discussed earlier.



**Figure 20.** Time-domain profiles of voltage and current during autonomous flight.

#### 4.7. Environmental Impact Assessment

To assess the environmental sustainability of the UAV system, a realistic estimate of its annual carbon footprint was made, considering energy usage, mission frequency, and the contribution of solar energy. Assuming the drone averages eight delivery missions per day, with each mission requiring approximately 88.8 Wh of energy, the total annual energy consumption is about 259 kWh. This amount reflects the total energy needed to complete 2920 missions per year (8 missions per day  $\times$  365 days). The drone is equipped with perovskite solar cells that allow for energy collection during the day. Assuming a 20% contribution from sunlight, the net energy drawn from the electrical grid decreases to approximately 207.2 kWh per year. Using the average European grid emission factor of 0.4 kg of CO<sub>2</sub> per kWh, the estimated annual emissions amount to 82.9 kg of CO<sub>2</sub>. This value is much lower than the emissions from traditional delivery vehicles. For comparison, a diesel van making the same number of deliveries would produce approximately 1000 kg of CO<sub>2</sub> annually, while an electric van powered by non-renewable grid energy would generate about 365 kg of CO<sub>2</sub> annually. All these estimates are based on equal annual energy use and standardised emission factors. These results confirm that the proposed UAV system, especially when powered by solar energy, has a significantly lower environmental impact. Its use in urban logistics could greatly help achieve emission reduction goals, particularly in areas where electric vehicle infrastructure is limited or underdeveloped.

To address the full environmental impact of the UAV system, a Life Cycle Assessment (LCA) was conducted, including production and end-of-life phases. The manufacturing stage, which involves carbon-reinforced PLA, brushless motors, LiPo batteries, and embedded electronics, is estimated to emit approximately 45.2 kg of CO<sub>2</sub> per unit, based on standard emission factors for composite materials and electronic components. The end-of-life phase, including disassembly, recycling of metals, and disposal of batteries and polymers, contributes an additional 12.6 kg of CO<sub>2</sub>. When combined with the operational emissions (82.9 kg CO<sub>2</sub>/year), the total annualised impact over a 5-year UAV lifespan is approximately 100.7 kg CO<sub>2</sub>/year. In contrast, a diesel delivery van emits around 1000 kg CO<sub>2</sub>/year, and an electric van powered by non-renewable grid energy emits about 365 kg CO<sub>2</sub>/year. These results confirm that the proposed UAV system offers a significantly lower environmental footprint across its entire lifecycle, reinforcing its suitability for sustainable urban logistics.

#### 4.8. System Integration and Urban Context

The integration of AI-powered navigation, solar-assisted energy systems, and light-weight structural materials has created a UAV platform capable of efficient, low-emission deliveries. The prototype serves as a base model for a startup focused on sustainable logistics, offering a scalable and cost-effective alternative to traditional delivery methods.

The statement that the prototype ‘serves as a basic model for a start-up’ is therefore indicative of a potential application scenario. Therefore, it does not imply that commercial implementation is immediate, but rather summarises a prospective path subject to the progressive experimental validation described in Section 3.10.

Along with its structural and energy innovations, the UAV uses a predictive AI system that improves mission planning and energy use. By using real-time environmental data and onboard inference, the system can choose energy-efficient flight paths, avoid risky areas, and adapt to changing wind conditions. The level of autonomy contributes to increased safety and durability, as well as reducing the environmental impact in crowded urban areas. The integration of embedded intelligence, renewable energy harvesting, and sustainable materials makes the UAV suitable for last-mile delivery services in smart cities. The modular design and minimal reliance on external infrastructure make the system suitable for use in areas with limited access to electric vehicle charging networks, reinforcing its role as a catalyst for green urban logistics.

#### 4.9. Inference Metrics and Decision Analysis

To evaluate the real-time usability of the DQN model on embedded hardware, we tested inference performance on the NVIDIA Jetson Nano using Tensor-Flow Lite (TFLite) with XNNPACK acceleration. The goal is to verify the model’s ability to make quick and accurate decisions within the constraints of edge computing. Empirical results showed that the average inference time per decision was approximately 3.2 milliseconds, enabling real-time control at over 300 decisions per second. This performance greatly surpasses the threshold needed for UAV navigation. In reinforcement learning, Equation (18) indicates the information an agent can receive by taking action on in-states and then following a specific policy  $\pi$ .

$$Q(s, a) = E \left[ \sum_{t=0}^{\infty} \gamma^t r_t | s_0 = s, a_0 = a \right] \quad (18)$$

where  $\gamma \in [0,1]$  is the discount factor and  $rt$  is the reward at time  $t$ . The DQN algorithm uses a neural network to approximate this function, up-dating the weights through the Q-learning method. To quantify the computational load of the trained models, we computed the number of multiply–accumulate operations (MACs) and floating-point operations (FLOPs) per inference. For a sequence of dense layers with sizes  $n_0, n_1, \dots, n_L$ , a forward pass requires approximately (19)

$$MAC_s = \sum_{i=0}^{L-1} n_i \cdot n_{i+1}, \quad \text{and} \quad FLOP_s \approx 2 \times MAC_s \quad (19)$$

For example, in the predictive model with input dimension 6 and layers [1,16,32], we obtain 720 MACs,  $\approx 1440$  FLOPs, and 769 trainable parameters.

This simple calculation illustrates the order of magnitude of the inference cost. Table 7 below summarizes the Q-values associated with each possible action in a representative inference scenario.

**Table 7.** Q-Value.

Action	Q-Value	Interpretation
Up	407.45	Good action, but not optimal
Down	345.91	Penalized, likely due to environmental risk
Left	351.82	Suboptimal
Right	399.26	Valid, but less efficient
Up-Right	420.33	Best choice—optimal trade-off
Down-Left	335.20	Penalized—high risk or inefficiency


We also measured the latency distributions on the Jetson Nano. The average inference time for the DQN agent was 3.2 ms ( $p_{95} \approx 5.1$  ms), while simultaneous execution with the EKF and image processing tasks resulted in an average of  $\approx 9.8$  ms ( $p_{95} \approx 14.6$  ms). These empirical results confirm that the model remains capable of operating in real-time, even with the increased dimensionality of the 3D state space, provided the inference latency remains below the control cycle deadline.

When deadlines are missed, techniques such as INT8 quantisation, pruning, or hardware offloading can be applied.

The command guides the drone towards the target, reducing both risk and energy consumption. The Q value of 420.33 indicates the total expected reward. This high value demonstrates the model's ability to prioritise safe and efficient routes.

The results determine how the DQN model performs accurate and fast inference on the Jetson Nano and how the TFLite + XNNPACK combination enables low-latency decision-making. The agent dynamically assesses the position and distance from the door, wind conditions, payload weight, and environmental risks. Figure 20 shows the Q-values for each action inferred by the DQN model on the Jetson Nano.

The action "Up-Right" has the highest Q-value of 420.33, confirming it as the optimal choice in the given scenario (Figure 21. Q-values for each action.

**Figure 21.** Q-values for each action.


The chart visually demonstrates the decision-making ability of the embedded AI system.

#### 4.10. Learning and Performance Metrics of the DQN

Various quantitative metrics, both during training and inference, determine the effectiveness of the DQN algorithm on the drone. Figures 20 and 21 illustrate the evolution of the metrics, providing a clear view of the learning process, decision accuracy, and the system's ability to adapt to complex urban contexts. To further validate stability, we assessed Lyapunov-inspired criteria: the decrease of temporal-difference error across episodes,

together with bounded Q-value variance, confirms convergence. These indicators complement success-rate metrics, showing that the closed-loop UAV + DQN system remains stable in the tested scenarios.

Specifically, Figure 22 illustrates the evolution of the parameter  $\epsilon$  (epsilon), which controls the balance between exploration and exploitation. Starting from 1.0,  $\epsilon$  decreases exponentially to 0.05, prompting the agent to transition from random exploration to using learnt policies. This change demonstrates the agent's growing confidence in their decision-making strategy.



**Figure 22.** Epsilon evolution in DQN model.

To validate that the observed convergence trends are not artefacts of a single run, we repeated the experiments with ten random seeds for each algorithmic variant. Table 8 reports the mean and standard deviation of cumulative reward and success rate, confirming that the growth patterns shown in Figures 15, 22 and 23 are statistically consistent.

**Table 8.** Statistical results over 10 random seeds for each algorithm.

Algorithm	Cumulative Reward (Mean $\pm$ Std)	Success Rate (Mean $\pm$ Std)
Vanilla DQN	142.3 $\pm$ 11.8	0.81 $\pm$ 0.05
Double DQN	156.7 $\pm$ 9.4	0.87 $\pm$ 0.04
Dueling DQNT	162.1 $\pm$ 7.6	0.89 $\pm$ 0.03


Figure 23 shows the success rate over the training episodes, defined as the percentage of simulations in which the UAV successfully reaches its target.

The curve consistently rises from approximately 10% to over 90%, demonstrating that the agent gradually learns more effective and stable navigation strategies.

The trends illustrated and the calculated values provide an improvement in policies and convergence. The decreasing  $\epsilon$  ensures that the agent depends more on its learned policy, while the rising success rate confirms that this policy results in consistently successful outcomes.

This behaviour shows that the DQN effectively adapts to the simulated urban environment, learning to avoid high-risk zones and optimise energy-efficient trajectories.

To finish analysing the learning metrics, we introduce the reward function, a crucial element in optimising the decision-making policies adopted by the agent.



**Figure 23.** Trend of the Operational Success of the DQN.

The reward function, essential in the reinforcement learning process, specifies how the agent evaluates the quality of its actions within the simulated environment. The function, designed to encourage behaviours that improve energy efficiency and operational safety, penalises exposure to high-risk environmental areas and the use of suboptimal routes. The reward structure allows for the shaping of exploratory behaviour that favours short and safe paths, discouraging entry into high-risk areas. For example, entering a cell with a risk level of 0.7 results in a reward of  $-4.5$ , prompting the agent to avoid such configurations in future training episodes. The trend of the reward function during training is shown in Figure 8, demonstrating a steady convergence toward optimal decision-making policies. The increasing average reward per episode and success rate highlight the effectiveness of the shaping strategy employed, confirming the model's ability to learn strong and adaptable behaviours in complex urban environments. Overall, the joint analysis of the  $\epsilon$  parameter, success rate, and reward function indicates a progressively more efficient and stable learning process, demonstrating the model's ability to adapt to complex scenarios and converge on optimal navigation strategies.

#### 4.11. Additional Evaluation: Impact of Action Discretization

The comparative simulation, testing three levels of action discretisation: Coarse (4 actions), Medium (8 actions), and Fine (16 actions), evaluates the limitations of DQN in continuous control tasks. Each configuration, tested over 100 episodes in a  $15 \times 15$  grid environment, measures navigation accuracy (% deviation from the optimal path), energy consumption (Wh), and mission success rate (% of successful episodes). Figure 23 and Table 9 summarise the results obtained. The paper concludes that finer action discretisation generally increases navigation accuracy and success rates of missions, albeit with diminishing returns beyond eight actions. Finer granularity also increases computational complexity and training times, suggesting a trade-off between accuracy of control and efficiency of the system.

The findings highlight the importance of developing more advanced reinforcement learning frameworks for continuous control in UAV missions.


**Table 9.** Performance.

Discretization Level	Navigation Accuracy (%)	Energy Consumption (Wh)	Mission Success Rate (%)
Coarse (4 actions)	~84	~96	~62
Medium (8 actions)	~88	~92	~75
Fine (16 actions)	~92	~88	~89

#### 4.11.1. Navigation Accuracy by Discretisation Level

Figure 24a shows the distribution of the drone's navigation accuracy for each level of action discretisation. Specifically, the following behaviours are obtained for the three types analysed:

- Coarse (4 actions): The drone has limited precision, with trajectories that are less close to optimal;
- Medium (8 actions): A significant improvement is observed, with smoother trajectories;
- Fine (16 actions): The drone achieves maximum precision thanks to its ability to perform more detailed manoeuvres.

**Figure 24.** Comparative performance of DQN under different action discretization levels.

This progressive improvement confirms that increasing the action discretisation enhances the agent's capacity to approximate continuous control, thereby enabling more accurate trajectory optimisation and reducing cumulative navigation errors.

#### 4.11.2. Energy Consumption per Discretisation Level

Figure 24b shows the average energy consumption for each level of discretisation:

- Coarse: Higher consumption due to less optimised manoeuvres;
- Medium: intermediate consumption, with greater efficiency;
- Fine: Lower fuel consumption, thanks to more direct trajectories and fewer corrections.

These results highlight the direct relationship between the spatial resolution of discretisation and the optimisation of UAV trajectories, confirming that finer discretisation levels enable more efficient path planning by reducing redundant manoeuvres and unnecessary energy expenditure.

#### 4.11.3. Mission Success Rate

The third diagram shows the percentage of episodes in which the drone completed the mission:

- Coarse: ~62%;
- Medium: ~75%;
- Fine: ~89%.

These results demonstrate that a finer discretisation of actions not only increases mission success rates but also improves the stability of the learning process, as the agent can generalise more effectively across complex scenarios. To avoid purely qualitative definitions of discretisation levels (e.g., “limited precision” or “smoother trajectories”), the action granularity is explicitly defined in terms of angular resolution and translational step size. In the proposed framework, the drone’s navigation space is discretised by setting the number of possible direction changes for each decision step.

Coarse discretisation corresponds to only four available actions (forward, 90° left, 90° right, and stop), medium discretisation allows for eight possible directions with 45° angular steps, and fine discretisation includes sixteen actions with 22.5° angular steps. In all cases, the translation unit step is normalised to one grid cell, while speed variations are introduced by allowing accelerations or decelerations within  $\pm 0.5$  m/s around the nominal cruising speed of 5 m/s.

Table 9 summarises the performance, while Table 10 summarises the relationship between the discretisation level, angular resolution, and resulting energy consumption. These quantitative definitions ensure that the impact of discretisation can be objectively evaluated in terms of navigation accuracy, energy cost, and trajectory smoothness.

**Table 10.** Quantitative definition of discretization levels and associated performance impact.

Discretization Level	Number Actions	Angular Resolution (°)	Max Heading Change per Step (°)	Velocity Step (m/s)	Average Energy Consumption (Wh/km)	Trajectory Deviation (m)
Coarse	4	90	$\pm 90$	$\pm 0.5$	14.8	2.4
Medium	8	45	$\pm 45$	$\pm 0.5$	12.6	1.5
Fine	16	22.5	$\pm 22.5$	$\pm 0.5$	11.9	0.9

The results demonstrate that finer discretisation significantly improves flight accuracy and energy efficiency, leading to higher mission success rates. However, they also highlight the inherent limitation of DQN in handling continuous control tasks such as pitch, roll, and thrust modulation. The results show that finer discretization improves trajectory adherence and reduces average deviation from the optimal path, at the cost of slightly higher computational load during training. Energy consumption decreases with higher action resolution, since smoother trajectories reduce redundant manoeuvres and sharp turns. These findings quantitatively support the choice of Fine discretization in scenarios where navigation accuracy and energy efficiency are critical.

#### 4.12. Simulation Environment Details

In pursuit of preserving the reproducibility and scientific integrity of the research outcomes presented herein, it is fundamental to provide an exhaustive description of the simulation environment, including software, hardware, and algorithmic details. All simulations were carried out with MATLAB/Simulink R2023b, and it was the development base of the UAV dynamics, energy management systems, as well as the environmental constructs. The reinforcement learning algorithms, including the therein proposed DQN policy, were defined with TensorFlow 2.12 in combination with Python 3.10 and integrated using MATLAB’s Reinforcement Learning Toolbox with the aim of preserving stable interoperability between the modules. The Deep Q-Network (DQN) agent training was carried out using a high-performance workstation with an Intel Core i7-12700K CPU, which is configured with 12 cores and 20 threads and a clock speed of 3.6 GHz, plus 32 GB RAM and an NVIDIA RTX 3080 GPU with 10 GB VRAM. The Intel Core i7-12700K CPU is manufactured by Intel Corporation, Santa Clara, CA, USA. The advanced computational hardware provided

fast training procedures, particularly in the handling of large replay buffers and parallelized simulation of the environments. After policy convergence, the optimised model was deployed and evaluated on an embeddable platform, namely the NVIDIA Jetson Nano, with a quad-core ARM Cortex-A57 CPU, 4 GB of LPDDR4 RAM, and a 128-core Maxwell GPU, and was used as the onboard computing system for the UAV demonstrator. The selection of the Jetson Nano was based on its desirable balance between processing and energy efficiency and thus served as a realistic constraint for environmentally friendly UAV missions. The simulation scenario used a  $15 \times 15$  discrete urban grid as the ground truth 2D scenario, later generalized to a 3D system with added height, vertical velocity, and turbulence dynamics. Important DQN training hyperparameters were optimised with care for policy stability and convergence: the learning rate was 0.001, the discount factor ( $\gamma$ ) was 0.95, the size of the experience replay buffer was 100,000 transitions, and the batch size was 64 samples. Exploration–exploitation trade-off was controlled with an  $\epsilon$ -greedy policy, with  $\epsilon$  decaying linearly from 1.0 down to 0.05 in 50,000 training iterations. The target network was updated every 500 iterations, and the gradient updates used the Adam optimiser. From a computational standpoint, the inferencing on the Jetson Nano was tested while running the DQN planner, sensor fusion (including the GNSS, IMU, and LiDAR), and energy estimation functions simultaneously. Average latency of the inferencing was found to be 47 milliseconds per decision step, equivalent to an actual control frequency of about 21 Hz and so capable of meeting specifications for real-time UAV control. Utilization of the GPU never exceeded as much as 65% even with peak multitasking usage, so there is some very minimal headroom available for other onboard functions, such as image processing enhancement or swarm command. By unveiling these specifications, the recent work allows reproducibility and transparency of the proposed eco-efficient UAV navigation system. Moreover, clear benchmarking of training and inference platforms allows for future implementation benchmarking and comparative analysis with other reinforcement learning methods applied in similar UAV scenarios.

## 5. Discussion

The proposed integrated system improved the operational efficiency and environmental sustainability of the UAV prototype. Adopting DQN allows the system to autonomously learn optimised navigation strategies. Compared to traditional route planning methods, the developed model offers superior adaptability and continuous performance improvement through iterative learning. The use of the NVIDIA Jetson Nano platform facilitates real-time processing of sensor data, including inputs from LiDAR, GPS, and inertial measurement units. This integrated processing solution supports deep learning models for both navigation and image enhancement, contributing to the drone's autonomous capabilities. The incorporation of convolutional neural networks (CNNs) and denoising autoencoders (AEs) further enhances the quality of the visual data, which is essential for accurately identifying and tracking delivery points. The results demonstrate a consistent development in the agent's behaviour, with an increase in the average reward and a success rate exceeding 90% in the final episodes. A controlled reduction in the epsilon parameter maintained an appropriate balance between exploration and exploitation, preventing premature convergence to suboptimal policies. Real-time inference on the Jetson Nano embedded platform showed that autonomous control without the need for cloud infrastructure was feasible from a computing standpoint. Consuming 16-bit numbers (FP16), a format that speeds up AI processes while consuming less energy, the Jetson Nano can do up to 472 billion operations per second. This feature guarantees real-time responsiveness by enabling the drone to make judgements in less than 4 milliseconds. Effective model convergence was also demonstrated by the training diagram analysis, which revealed a steady decrease in the loss function.

The system's responsiveness was further enhanced by dynamic decision speed adaptation, which qualified it for missions in real-world settings. The methodology used in this work is rigorous and multidisciplinary, despite the fact that it is dependent on numerical simulations. Compared to state-of-the-art UAVs for urban logistics, the proposed prototype demonstrates an optimal trade-off between payload capacity, endurance, and eco-efficiency, aligning with the goals of the European Green Deal. The simulation framework integrates CAD-based aerodynamic modelling, energy consumption simulations based on real BLDC motor and LiPo battery discharge profiles, and the training of a Deep Q-Network (DQN) in a structured virtual environment. The system's predictive performance, validated through a comparison between estimated and simulated energy consumption, has a mean absolute error (MAE) of 15.37 Wh and a mean squared error (MSE) of 375.56 Wh. The results confirm the reliability of the proposed architecture and its suitability for real-time applications. The absence of physical testing at this stage is a deliberate methodological choice, aimed at optimising computational resources and simulation accuracy. By lowering external variables and guaranteeing consistency across test situations, this method enables a thorough assessment of control algorithms and the energy architecture in a virtual environment. Additionally, it facilitates the iterative development of AI-based decision-making techniques and aids in the early detection of design bottlenecks. All things considered, the created prototype offers a viable way to handle sustainable urban logistics, particularly in places with inadequate infrastructure for electric vehicles. A scalable, effective, and eco-friendly platform is provided by the integration of biodegradable materials, solar energy, and artificial intelligence. This creates real prospects for commercial applications and creative businesses in the smart mobility space. Nevertheless, the UAV system has the following serious problems. At this stage, the drone prototype has not yet been built, and all evaluations were performed within a simulated  $15 \times 15$  urban environment. While simulation results are promising, several practical deployment challenges remain before operational use: (i) Real-world variability and domain transfer—sensor noise, GNSS outages (urban canyons), and unmodelled aerodynamics may reduce transferability; field testing with Hardware-in-the-Loop (HIL) and phased outdoor trials are required. (ii) Airspace integration and regulation—compliance with EASA/FAA UAS rules, geofencing, integration with U-space/ATC services, and certification paths for BVLOS operations must be addressed. (iii) Safety and detect-and-avoid—formal detect-and-avoid (DAA) and redundancy (GNSS fallback, vision/LiDAR) strategies are necessary for certification. (iv) PV durability and environmental exposure—perovskite PV cells require encapsulation and testing for humidity/thermal cycling prior to operational deployment. This methodological choice ensured the controlled validation of the energy architecture and the DQN-based navigation system. It is worth noting that standard DQN does not provide formal Lyapunov-type guarantees of stability or convergence in closed-loop UAV control. Our analysis is therefore empirical and based on simulation tests of convergence and policy robustness. Future research will explore Safe Reinforcement Learning and Lyapunov-based critical networks, barrier function methods, to ensure constraint satisfaction and long-term stability in real-world implementation. Future studies will focus on building the prototype and conducting regulated flight tests to confirm the results. The proposed study is part of a multi-phase validation process, necessary to strengthen methodological transparency and to highlight the feasibility of subsequent experiments. Although the system integrates sensors for wind and hazard detection, its performance in extreme weather conditions still needs to be experimentally validated. The efficiency of the Jetson Nano imposes constraints on the complexity of AI models that can be run in real time. More sophisticated algorithms might require higher-performance hardware. To address the identified limitations, test campaigns need to be conducted in controlled urban environments to evaluate the system's robustness

in the presence of air traffic, dynamic obstacles, and electromagnetic interference. These experiments will allow us to calibrate the navigation models and refine the decision-making strategies learnt in simulation. The incorporation of advanced DQN variants, as well as the adoption of multi-agent approaches, can significantly improve the system's ability to operate in complex and dynamic environments. Future studies might look into hybrid architectures that combine DQN with continuous control models for low-level actuation. These architectures allow for a richer representation of decision-making policies because they contextualise the performance of the DQN model within the broader landscape of reinforcement learning algorithms. The performance of the DQN should be compared to both the most advanced and the less capable reinforcement learning algorithms. The TD3 and DDPG control models demonstrated superior performance in terms of navigation accuracy, energy efficiency, and mission success rate. However, these improvements come with significantly higher computational complexity, which limits their deployability on embedded platforms such as the Jetson Nano. Their actor–critic architectures demand more memory, dual-network inference, and frequent updates, which can hinder real-time responsiveness and energy autonomy—two critical factors in UAV operations. However, in organised navigation tasks, DQN has been demonstrated to perform better than more straightforward models like SARSA and A2C. While A2C, despite its theoretical advantages, requires significant tuning and processing resources to match DQN's stability, SARSA frequently encounters convergence issues and has trouble generalising in dynamic contexts. In comparative studies, DQN was the only algorithm to complete navigation tasks in simulated urban grids reliably, demonstrating its robustness and dependability in discrete action spaces. This dual positioning emphasizes the reason for our choice: DQN provides a practical balance between performance, simplicity, and hardware compatibility. It is especially suitable for high-level path planning in embedded UAV systems, where computational limits and real-time decision-making are critical. Future work might look into hybrid architectures that combine DQN with continuous control models for low-level actuation. Still, the current approach focuses on feasibility and robustness in real-world deployment scenarios.

## 6. Conclusions

This study demonstrates the feasibility of combining embedded artificial intelligence, renewable energy harvesting, and eco-efficient design to enable autonomous and sustainable drone operations for urban logistics. Within a simulation environment, the proposed system integrates a decision-making algorithm based on DQN running on low-power hardware with perovskite-assisted solar harvesting and an optimised propulsion-storage chain. The results highlight performance benefits in representative urban scenarios, including increased range and improved energy efficiency. The study is based solely on simulation results, without physical prototyping or flight testing. Aerodynamic effects, environmental variability, soiling and shading of photovoltaic panels, and onboard computational limitations can significantly influence real-world performance. The DQN policy was trained under simplified stochastic disturbance conditions, and its ability to generalise to unknown environments has yet to be validated. Future steps will focus on building a flight-ready prototype equipped with perovskite photovoltaic modules and a real-time MPPT controller, followed by controlled test campaigns. Further objectives include characterising photovoltaic yield and MPPT dynamics under flight conditions, evaluating safety and reliability through hardware-in-the-loop experiments, and extending the reinforcement learning framework towards multi-objective optimisation covering energy constraints, safety margins, and level of service. Future developments will focus on real-world prototyping and regulated test campaigns, enabling validation of the proposed eco-efficient

UAV within operational urban logistics networks. In addition to these prospects, a gradual prototype validation plan has been defined. The first phase involves hardware-in-the-loop (HIL) experiments, in which data streams from GNSS, IMU and LiDAR sensors are fed into the EKF-DQN pipeline with real-time constraints. The second phase involves laboratory validation with perovskite photovoltaic modules, brushless DC motors and Li-Po batteries, focusing on MPPT-BMS coordination and energy distribution under controlled irradiation and load conditions. Finally, incremental outdoor campaigns will be conducted to progressively test navigation accuracy, energy yield and robustness against turbulence and environmental variability. This roadmap ensures engineering feasibility and bridges the gap between simulation-only testing and field implementation.

Beyond its technical contributions, the proposed system provides a replicable design paradigm that can inspire both academic research and industrial innovation in sustainable logistics. Its scalability, modularity, and compatibility with decentralised infrastructure make it a viable candidate for integration into smart city ecosystems, thereby reinforcing its broader socio-economic and environmental relevance.

**Author Contributions:** Conceptualization, L.B., G.A. and F.C.; methodology, L.B., F.C.; software, L.B., G.A., F.C.; validation, L.B., G.B., F.L.; formal analysis, L.B., F.L., G.M.M., G.A., G.B. and F.C.; investigation, L.B., G.A.; data curation, L.B., G.B.; writing original draft preparation, L.B., F.L., G.B., G.A., F.C.; writing; supervision, L.B., F.L., G.A. All authors have read and agreed to the published version of the manuscript.

**Funding:** This research received no external funding.

**Data Availability Statement:** The original contributions presented in this study are included in the article. Further inquiries can be directed to the corresponding authors.

**Conflicts of Interest:** Author Francesco Cotroneo was employed by the company Nophys srl. The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

## Abbreviations

UAV	Unmanned Aerial Vehicle
AI	Artificial Intelligence
DQN	Deep Q-Network
BLDC	Brushless DC Motor
LiPo	Lithium Polymer (Battery)
IMU	Inertial Measurement Unit
LiDAR	Light Detection and Ranging
DRL	Deep reinforcement learning
GNSS	Global Navigation Satellite System
EGNOS	European Geostationary Navigation Overlay Service
INS	Inertial Navigation System
EKF	Extended Kalman Filter
CNN	Convolutional Neural Network
EDSR	Enhanced Deep Super-Resolution
DAA	Detect-and-avoid
DAE	Denoising Autoencoder
DDPG	Deep Deterministic Policy Gradient
PF	Particle Filtering
VIO	Visual-Inertial Odometry

BMS	Battery Management System
TD3	Twin Delayed DDPG
UKF	Unscented Kalman Filter
RRT	Rapidly Random Trees
MPPT	Maximum Power Point Tracking
PWM	Pulse Width Modulation
ESC	Electronic Speed Controller
PWR	Payload-to-Weight Ratio
L/D	Lift-to-Drag Ratio
FLOPS	Floating Point Operations Per Second
TFLite	TensorFlow Lite
RIS	Reconfigurable Intelligent Surfaces
MBSE	Model-Based Systems Engineering

## References

- Jahani, H.; Khosravi, Y.; Kargar, B.; Ong, K.L.; Arisian, S. Exploring the role of drones and UAVs in logistics and supply chain management: A novel text-based literature review. *Int. J. Prod. Res.* **2025**, *63*, 1873–1897. [\[CrossRef\]](#)
- Telli, K.; Kraa, O.; Himeur, Y.; Ouamane, A.; Boumechraz, M.; Atalla, S.; Mansoor, W. A Comprehensive Review of Recent Research Trends on Unmanned Aerial Vehicles (UAVs). *Systems* **2023**, *11*, 400. [\[CrossRef\]](#)
- Betti Sorbelli, F. UAV-based delivery systems: A systematic review, current trends, and research challenges. *J. Auton. Transp. Syst.* **2024**, *1*, 1–40. [\[CrossRef\]](#)
- Mohsan, S.A.H.; Othman, N.Q.H.; Li, Y.; Alsharif, M.H.; Khan, M.A. Unmanned aerial vehicles (UAVs): Practical aspects, applications, open challenges, security issues, and future trends. *Intell. Serv. Robot.* **2023**, *16*, 109–137. [\[CrossRef\]](#)
- Rejeb, A.; Rejeb, K.; Simske, S.J.; Treiblmaier, H. Drones for supply chain management and logistics: A review and research agenda. *Int. J. Logist. Res. Appl.* **2023**, *26*, 708–731. [\[CrossRef\]](#)
- Kumar, A.; Prybutok, V.; Sangana, V.K.R. Environmental Implications of Drone-Based Delivery Systems: A Structured Literature Review. *Clean Technol.* **2025**, *7*, 24. [\[CrossRef\]](#)
- Mohanraj, D.; ArulDavid, R.; Verma, R.; Sathiyasekar, K.; Barnawi, A.B.; Chokkalingam, B.; Mihet-Popa, L. A review of BLDC motor: State of art, advanced control techniques, and applications. *IEEE Access* **2022**, *10*, 54833–54869. [\[CrossRef\]](#)
- Yaşa, Y. An Efficient Brushless DC Motor Design for Unmanned Aerial Vehicles. *Eur. J. Sci. Technol.* **2022**, *35*, 288–294. [\[CrossRef\]](#)
- Emimi, M.; Khaleel, M.; Alkrash, A. The Current Opportunities and Challenges in Drone Technology. *Int. J. Electr. Eng. Sustain.* **2023**, *1*, 74–89.
- Senhaji-Mouhaddib, I.; Bouzakri, S.; Lagrat, I. Advanced autonomous navigation technologies for uavs: Challenges and opportunities. In Proceedings of the 2024 4th International Conference on Innovative Research in Applied Science, Engineering and Technology, Fez, Morocco, 16–17 May 2024.
- Laganà, F.; Facci, A.R. Parametric optimisation of a pulmonary ventilator using the Taguchi method. *J. Electr. Eng.* **2025**, *76*, 265–274. [\[CrossRef\]](#)
- Khurshid, N.; Kumar, R.; Vishwakarma, D.K.; Kumar, S.; Pandit, B.A.; Khan, I.; Pukhta MSYadav, K.K. Interaction Modeling of Surface Water and Groundwater: An Evaluation of Current and Future Issues. *Water Conserv. Sci. Eng.* **2025**, *10*, 37. [\[CrossRef\]](#)
- Laganà, F.; Pellicanò, D.; Arruzzo, M.; Praticò, D.; Pullano, S.A.; Fiorillo, A.S. FEM-Based Modelling and AI-Enhanced Monitoring System for Upper Limb Rehabilitation. *Electronics* **2025**, *14*, 2268. [\[CrossRef\]](#)
- Praticò, D.; Laganà, F.; Oliva, G.; Fiorillo, A.S.; Pullano, S.A.; Calcagno, S.; De Carlo, D.; La Foresta, F. Sensors and Integrated Electronic Circuits for Monitoring Machinery On Wastewater Treatment: Artificial Intelligence Approach. In Proceedings of the 2024 IEEE Sensors Applications Symposium (SAS), Naples, Italy, 23–25 July 2024; pp. 1–6. [\[CrossRef\]](#)
- Praticò, D.; Laganà, F.; Oliva, G.; Fiorillo, A.S.; Pullano, S.A.; Calcagno, S.; De Carlo, D.; La Foresta, F. Integration of LSTM and U-Net Models for Monitoring Electrical Absorption With a System of Sensors and Electronic Circuits. *IEEE Trans. Instrum. Meas.* **2025**, *74*, 2533311. [\[CrossRef\]](#)
- Menniti, M.; Oliva, G.; Laganà, F.; Bianco, M.G.; Fiorillo, A.S.; Pullano, S.A. Portable Non-Invasive Ventilator for Homecare and Patients Monitoring System. In Proceedings of the 2023 IEEE International Symposium on Medical Measurements and Applications (MeMeA), Jeju, Republic of Korea, 26 May 2023; pp. 1–5. [\[CrossRef\]](#)
- IEC 62133-2:2017; Secondary Cells and Batteries Containing Alkaline or Other Non-Acid Electrolytes—Safety Requirements for Portable Sealed Secondary Cells, and for Batteries Made from Them, for Use in Portable Applications—Part 2: Lithium Systems. International Electrotechnical Commission: Geneva, Switzerland, 2017.

18. Massaoudi, M.S.; Abu-Rub, H.; Ghayeb, A. Navigating the landscape of deep reinforcement learning for power system stability control: A review. *IEEE Access* **2023**, *11*, 134298–134317. [[CrossRef](#)]
19. Li, S.; Fang, Z.; Verma, S.C.; Wei, J.; Savkin, A.V. Navigation and Deployment of Solar-Powered Unmanned Aerial Vehicles for Civilian Applications: A Comprehensive Review. *Drones* **2024**, *8*, 42. [[CrossRef](#)]
20. Saravanakumar, Y.N.; Sultan, M.T.H.; Shahar, F.S.; Giernacki, W.; Łukaszewicz, A.; Nowakowski, M.; Holovatyy, A.; Stępień, S. Power Sources for Unmanned Aerial Vehicles: A State-of-the Art. *Appl. Sci.* **2023**, *13*, 11932. [[CrossRef](#)]
21. Wei, Z.; Wang, S.; Chen, K.; Wang, F. ROS-Based Navigation and Obstacle Avoidance: A Study of Architectures, Methods, and Trends. *Sensors* **2025**, *25*, 4306. [[CrossRef](#)]
22. Pastor, R.; Rodriguez, P.G.; Lecuona, A.; Cortés, J.P. Smart Agri-Region and Value Engineering. *Systems* **2025**, *13*, 430. [[CrossRef](#)]
23. AL-Dosari, K.; Fetais, N. A New Shift in Implementing Unmanned Aerial Vehicles (UAVs) in the Safety and Security of Smart Cities: A Systematic Literature Review. *Safety* **2023**, *9*, 64. [[CrossRef](#)]
24. Karatzas, A.; Karras, A.; Karras, C.; Giotopoulos, K.C.; Oikonomou, K.; Sioutas, S. On autonomous drone navigation using deep learning and an intelligent rainbow DQN agent. In Proceedings of the International Conference on Intelligent Data Engineering and Automated Learning, Manchester, UK, 24–26 November 2022; pp. 134–145.
25. Amarasooriya, P.M.D.S.; Sandaruwan, K.D. AI-Based 3D Simulation for Drone Flight Dynamics. *Int. J. Adv. ICT Emerg. Reg.* **2025**, *18*, 2. [[CrossRef](#)]
26. Puente-Castro, A.; Rivero, D.; Pazos, A.; Fernandez-Blanco, E. A review of artificial intelligence applied to path planning in UAV swarms. *Neural Comput. Appl.* **2022**, *34*, 153–170. [[CrossRef](#)]
27. Reda, M.; Onsy, A.; Haikal, A.Y.; Ghanbari, A. Path planning algorithms in the autonomous driving system: A comprehensive review. *Robot. Auton. Syst.* **2024**, *174*, 104630. [[CrossRef](#)]
28. Debnath, D.; Vanegas, F.; Sandino, J.; Hawary, A.F.; Gonzalez, F. A Review of UAV Path-Planning Algorithms and Obstacle Avoidance Methods for Remote Sensing Applications. *Remote Sens.* **2024**, *16*, 4019. [[CrossRef](#)]
29. Xu, J. Efficient trajectory optimization and resource allocation in UAV 5G networks using dueling-Deep-Q-Networks. *Wirel. Netw.* **2024**, *30*, 6687–6697. [[CrossRef](#)]
30. El-Hajj, M. Enhancing Communication Networks in the New Era with Artificial Intelligence: Techniques, Applications, and Future Directions. *Network* **2025**, *5*, 1. [[CrossRef](#)]
31. He, J.; Pan, C.; Wang, J.; Lu, C.; Chen, S. Federated Deep Q-network: A Dynamic Task Allocation Strategy for UAV-Assisted Cell-Free Networks. In Proceedings of the International Conference on Frontiers of Electronics, Information and Computation Technologies, Singapore, 10 May 2025; Springer Nature: Berlin, Germany, 2025; pp. 451–459.
32. Qin, H.; Shao, S.; Wang, T.; Yu, X.; Jiang, Y.; Cao, Z. Review of Autonomous Path Planning Algorithms for Mobile Robots. *Drones* **2023**, *7*, 211. [[CrossRef](#)]
33. Kabir, H.; Tham, M.L.; Chang, Y.C. Internet of robotic things for mobile robots: Concepts, technologies, challenges, applications, and future directions. *Digit. Commun. Netw.* **2023**, *9*, 1265–1290. [[CrossRef](#)]
34. Chang, Z.; Deng, H.; You, L.; Min, G.; Garg, S.; Kaddoum, G. Trajectory design and resource allocation for multi-UAV networks: Deep reinforcement learning approaches. *IEEE Trans. Netw. Sci. Eng.* **2022**, *10*, 2940–2951. [[CrossRef](#)]
35. Wang, T.; Huang, X.; Wu, Y.; Qian, L.; Lin, B.; Su, Z. UAV swarm-assisted two-tier hierarchical federated learning. *IEEE Trans. Netw. Sci. Eng.* **2023**, *11*, 943–956. [[CrossRef](#)]
36. Peng, C.; Wu, Z.; Huang, X.; Wu, Y.; Kang, J.; Huang, Q.; Xie, S. Joint energy and completion time difference minimization for UAV-enabled intelligent transportation systems: A constrained multi-objective optimization approach. *IEEE Trans. Intell. Transp. Syst.* **2024**, *25*, 14040–14053. [[CrossRef](#)]
37. Zhang, H.; Huang, M.; Zhou, H.; Wang, X.; Wang, N.; Long, K. Capacity maximization in RIS-UAV networks: A DDQN-based trajectory and phase shift optimization approach. *IEEE Trans. Wirel. Commun.* **2022**, *22*, 2583–2591. [[CrossRef](#)]
38. Liu, R.; Guo, K.; Li, X.; Dev, K.; Khowaja, S.A.; Tsiftsis, T.A.; Song, H. RIS-empowered satellite-aerial-terrestrial networks with PD-NOMA. *IEEE Commun. Surv. Tutor.* **2024**, *26*, 2258–2289. [[CrossRef](#)]
39. Rai, S.; Rawat, A.; Kumar, A. Design and performance analysis of high-altitude UAVs: Trends, challenges, and innovations. *Discov. Appl. Sci.* **2025**, *7*, 833. [[CrossRef](#)]
40. Zhang, Y.; Zhao, R.; Mishra, D.; Ng, D.W.K. A Comprehensive Review of Energy-Efficient Techniques for UAV-Assisted Industrial Wireless Networks. *Energies* **2024**, *17*, 4737. [[CrossRef](#)]
41. Liu, H.; Li, J.; Yang, H.; Wang, J.; Li, B.; Zhang, H.; Yi, Y. TiN-Only Metasurface Absorber for Solar Energy Harvesting. *Photonics* **2025**, *12*, 443. [[CrossRef](#)]
42. Noman, M.; Khan, Z.; Jan, S.T. A comprehensive review on the advancements and challenges in perovskite solar cell technology. *RSC Adv.* **2024**, *14*, 5085–5131. [[CrossRef](#)]
43. Meng, W.; Zhang, X.; Zhou, L.; Guo, H.; Hu, X. Advances in UAV Path Planning: A Comprehensive Review of Methods, Challenges, and Future Directions. *Drones* **2025**, *9*, 376. [[CrossRef](#)]

44. Tao, B.; Kim, J.H. Deep reinforcement learning-based local path planning in dynamic environments for mobile robot. *J. King Saud Univ.-Comput. Inf. Sci.* **2024**, *36*, 102254. [CrossRef]
45. Wanner, D.; Hashim, H.A.; Srivastava, S.; Steinhauer, A. UAV avionics safety, certification, accidents, redundancy, integrity, and reliability: A comprehensive review and future trends. *Drone Syst. Appl.* **2024**, *12*, 1–23. [CrossRef]
46. Aftab, S.; Hussain, S.; Kabir, F.; Aslam, M.; Rajpar, A.H.; Al-Sehemi, A.G. Advances in flexible perovskite solar cells: A comprehensive review. *Nano Energy* **2024**, *120*, 109112. [CrossRef]
47. Kutsarov, D.I.; Rezaee, E.; Lambert, J.; Stroud, W.T.; Panagiotopoulos, A.; Silva, S.R.P. Progress in Flexible Perovskite Solar Cells: Paving the Way for Scalable Manufacturing. *Adv. Mater. Technol.* **2025**, 2401834. [CrossRef]
48. Baidya, T.; Nabi, A.; Moh, S. Trajectory-Aware Offloading Decision in UAV-Aided Edge Computing: A Comprehensive Survey. *Sensors* **2024**, *24*, 1837. [CrossRef]
49. Shuaibu, A.S.; Mahmoud, A.S.; Sheltami, T.R. A Review of Last-Mile Delivery Optimization: Strategies, Technologies, Drone Integration, and Future Trends. *Drones* **2025**, *9*, 158. [CrossRef]
50. Yu, S.; Lee, H.; Ju, C.; Han, H. Enhanced DBR mirror design via D3QN: A reinforcement learning approach. *PLoS ONE* **2024**, *19*, e0307211. [CrossRef] [PubMed]
51. Huo, Y.; Gang, S.; Guan, C. FCIHMT: Feature Cross-Layer Interaction Hybrid Method Based on Res2Net and Transformer for Remote Sensing Scene Classification. *Electronics* **2023**, *12*, 4362. [CrossRef]
52. Al-Haddad, L.A.; Jaber, A.A.; Giernacki, W.; Khan, Z.H.; Ali, K.M.; Tawafik, M.A.; Humaidi, A.J. Quadcopter unmanned aerial vehicle structural design using an integrated approach of topology optimization and additive manufacturing. *Designs* **2024**, *8*, 58. [CrossRef]
53. Angiulli, G.; Calcagno, S.; De Carlo, D.; Laganá, F.; Versaci, M. Second-Order Parabolic Equation to Model, Analyze, and Forecast Thermal-Stress Distribution in Aircraft Plate Attack Wing–Fuselage. *Mathematics* **2020**, *8*, 6. [CrossRef]
54. Anand, S.; Mishra, A.K. High-performance materials used for UAV manufacturing: Classified review. *Int. J. All Res. Educ. Sci. Methods (IJARESM)* **2022**, *10*, 2455–6211.
55. DIN EN 60728-1:2015-03; Cable Networks for Television Signals, Sound Signals and Interactive Services—Part 1: System Performance of Forward Paths. Deutsches Institut für Normung: Berlin, Germany, 2015.
56. Advantages and Disadvantages of Brushless DC (BLDC) Motor—Electricalsblog. (s.d.). Available online: <https://arshon.com/blog/advantages-and-disadvantages-of-brushless-dc-motors-blDC/> (accessed on 30 January 2025).
57. Hernández-Guzmán, V.M.; Silva-Ortigoza, R.; Orrante-Sakanassi, J.A. Brushless DC-Motor. In *Energy-Based Control of Electromechanical Systems*; Springer: Berlin/Heidelberg, Germany, 2020. [CrossRef]
58. ISO 12345:2022; Diesel Engines—Cleanliness Assessment of Fuel Injection Equipment. International Organization for Standardization: Geneva, Switzerland, 2022.
59. Tattu. Tattu 6S 12000mAh 25C LiPo Battery: Technical Datasheet and Specifications. Gens Ace Tattu. 2023. Available online: <https://www.genstattu.com> (accessed on 30 August 2025).
60. El-Atab, N.; Mishra, R.B.; Alshambari, R.; Hussain, M.M. Solar powered small unmanned aerial vehicles: A review. *Energy Technol.* **2021**, *9*, 2100587. [CrossRef]
61. Sornek, K.; Augustyn-Nadzieja, J.; Rosikoń, I.; Łopusiewicz, R.; Łopusiewicz, M. Status and Development Prospects of Solar-Powered Unmanned Aerial Vehicles—A Literature Review. *Energies* **2025**, *18*, 1924. [CrossRef]
62. Jarraya, I.; Al-Batati, A.; Kadri, M.B.; Abdelkader, M.; Ammar, A.; Boulila, W.; Koubaa, A. Gnss-denied unmanned aerial vehicle navigation: Analyzing computational complexity, sensor fusion, and localization methodologies. *Satell. Navig.* **2025**, *6*, 9. [CrossRef]
63. Uzun, F.N.; Kayrici, M.; Akkuzu, B. *Nvidia Jetson Nano Development Kit*; ISRES Publishing: Singapore, 2021.
64. Diamantidou, E.; Lalas, A.; Votis, K.; Tzovaras, D. Multimodal deep learning framework for enhanced accuracy of UAV detection. In Proceedings of the International Conference on Computer Vision Systems, Thessaloniki, Greece, 23–25 September 2019; pp. 768–777.
65. Pal, O.K.; Shovon, M.S.H.; Mridha, M.F.; Shin, J. In-depth review of AI-enabled unmanned aerial vehicles: Trends, vision, and challenges. *Discov. Artif. Intell.* **2024**, *4*, 97. [CrossRef]
66. Kalenberg, K.; Müller, H.; Polonelli, T.; Schiaffino, A.; Niculescu, V.; Cioflan, C.; Benini, L. Stargate: Multimodal sensor fusion for autonomous navigation on miniaturised UAVs. *IEEE Internet Things J.* **2024**, *11*, 21372–21390. [CrossRef]
67. Hafiz, A.M. A survey of deep q-networks used for reinforcement learning: State of the art. In Proceedings of the Intelligent Communication Technologies and Virtual Mobile Networks: Proceedings of ICICV, Tirunelveli, India, 10–11 February 2022; pp. 393–402.
68. Sharma, G.; Jain, S.; Sharma, R.S. Path Planning for Fully Autonomous UAVs—A Taxonomic Review and Future Perspectives. *IEEE Access* **2025**, *13*, 13356–13379. [CrossRef]
69. Zhang, D.; Wang, Y.; Meng, L.; Yan, J.; Qin, C. Adaptive critic design for safety-optimal FTC of unknown nonlinear systems with asymmetric constrained-input. *ISA Trans.* **2024**, *155*, 309–318. [CrossRef]

70. Qin, C.; Hou, S.; Pang, M.; Wang, Z.; Zhang, D. Reinforcement learning-based secure tracking control for nonlinear interconnected systems: An event-triggered solution approach. *Eng. Appl. Artif. Intell.* **2025**, *161*, 112243. [CrossRef]
71. Qin, C.; Pang, M.; Wang, Z.; Hou, S.; Zhang, D. Observer based fault tolerant control design for saturated nonlinear systems with full state constraints via a novel event-triggered mechanism. *Eng. Appl. Artif. Intell.* **2025**, *161*, 112221. [CrossRef]
72. Qin, C.; Qiao, X.; Wang, J.; Zhang, D.; Hou, Y.; Hu, S. Barrier-critic adaptive robust control of nonzero-sum differential games for uncertain nonlinear systems with state constraints. *IEEE Trans. Syst. Man Cybern. Syst.* **2023**, *54*, 50–63. [CrossRef]
73. Hu, D.; Minner, J. UAVs and 3D City Modeling to Aid Urban Planning and Historic Preservation: A Systematic Review. *Remote Sens.* **2023**, *15*, 5507. [CrossRef]
74. Praticò, D.; Carlo, D.D.; Silipo, G.; Laganà, F. Hybrid FEM-AI Approach for Thermographic Monitoring of Biomedical Electronic Devices. *Computers* **2025**, *14*, 344. [CrossRef]
75. Praticò, D.; Laganà, F. Infrared Thermographic Signal Analysis of Bioactive Edible Oils Using CNNs for Quality Assessment. *Signals* **2025**, *6*, 38. [CrossRef]
76. Available online: [https://www.google.com/maps/dir/Aeroporto+dello+Stretto+%22Tito+Minniti%22+\(REG\),+Via+Ravagnese,+Reggio+Calabria,+RC/Vico+Friuli+I',+89124+Reggio+Calabria+RC/@38.0893991,15.6443371,10441m/data=!3m1!1e3!4m13!4m12!1m5!1m1!1s0x13145a81629f1beb:0xe533be0acb04f635!2m2!1d15.6538057!2d38.0739829!1m5!1m1!1s0x13145093cb4e2c0f:0x806dfced4a36f843!2m2!1d15.6559073!2d38.1146353?authuser=0&entry=tту&g\\_ep=EgoyMDI1MDkwMy4wIKXMDSoASAFQAw%3D%3D](https://www.google.com/maps/dir/Aeroporto+dello+Stretto+%22Tito+Minniti%22+(REG),+Via+Ravagnese,+Reggio+Calabria,+RC/Vico+Friuli+I',+89124+Reggio+Calabria+RC/@38.0893991,15.6443371,10441m/data=!3m1!1e3!4m13!4m12!1m5!1m1!1s0x13145a81629f1beb:0xe533be0acb04f635!2m2!1d15.6538057!2d38.0739829!1m5!1m1!1s0x13145093cb4e2c0f:0x806dfced4a36f843!2m2!1d15.6559073!2d38.1146353?authuser=0&entry=tту&g_ep=EgoyMDI1MDkwMy4wIKXMDSoASAFQAw%3D%3D) (accessed on 6 September 2025).
77. El Sakka, M.; Ivanovici, M.; Chaari, L.; Mothe, J. A Review of CNN Applications in Smart Agriculture Using Multimodal Data. *Sensors* **2025**, *25*, 472. [CrossRef]
78. Vedrtnam, A.; Negi, H.; Kalauni, K. Materials and Energy-Centric Life Cycle Assessment for Drones: A Review. *J. Compos. Sci.* **2025**, *9*, 169. [CrossRef]
79. Wei, H.; Lou, B.; Zhang, Z.; Liang, B.; Wang, F.Y.; Lv, C. Autonomous navigation for eVTOL: Review and future perspectives. *IEEE Trans. Intell. Veh.* **2024**, *9*, 4145–4171. [CrossRef]
80. Olthuis, J.J.; Sciancalepore, S.; Zannone, N. Cyberattacks and defenses for Autonomous Navigation Systems: A systematic literature review. *Comput. Netw.* **2025**, *267*, 111331. [CrossRef]
81. Laganà, F.; Praticò, D.; De Carlo, D.; Oliva, G.; Pullano, S.A.; Calcagno, S. Engineering Biomedical Problems to Detect Carcinomas: A Tomographic Impedance Approach. *Eng* **2024**, *5*, 1594–1614. [CrossRef]
82. Laganà, F.; De Carlo, D.; Calcagno, S.; Pullano, S.A.; Critello, D.; Falcone, F.; Fiorillo, A.S. Computational model of cell deformation under fluid flow based rolling. In Proceedings of the E-Health and Bioengineering Conference (EHB), Iasi, Romania, 21–23 November 2019; pp. 1–4. [CrossRef]
83. Laganà, F.; Pullano, S.A.; Angiulli, G.; Versaci, M. Optimized Analytical–Numerical Procedure for Ultrasonic Sludge Treatment for Agricultural Use. *Algorithms* **2024**, *17*, 592. [CrossRef]
84. Bibbò, L.; Angiulli, G.; Laganà, F.; Praticò, D.; Cotroneo, F.; La Foresta, F.; Versaci, M. MEMS and IoT in HAR: Effective Monitoring for the Health of Older People. *Appl. Sci.* **2025**, *15*, 4306. [CrossRef]

**Disclaimer/Publisher’s Note:** The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.